arXiv:1607.02649v1 [cs.IT] 9 Jul 2016

Linear signal recovery from b-bit-quantized linear measurements: precise analysis of the trade-off between bit depth and number of measurements

Martin Slawski
Department of Statistics and Biostatistics
Department of Computer Science
Rutgers University
Piscataway, NJ 08854, USA
martin.slawski@rutgers.edu

Ping Li
Department of Statistics and Biostatistics
Department of Computer Science
Rutgers University
Piscataway, NJ 08854, USA
pingli@stat.rutgers.edu

Abstract

We consider the problem of recovering a high-dimensional structured signal from independent Gaussian linear measurements each of which is quantized to b bits. Our interest is in linear approaches to signal recovery, where "linear" means that non-linearity resulting from quantization is ignored and the observations are treated as if they arose from a linear measurement model. Specifically, the focus is on a generalization of a method for one-bit observations due to Plan and Vershynin [*IEEE Trans. Inform. Theory,* ***59*** *(2013), 482–494*]. At the heart of the present paper is a precise characterization of the optimal trade-off between the number of measurements m and the bit depth per measurement b given a total budget of $B = m \cdot b$ bits when the goal is to minimize the ℓ_2-error in estimating the signal. It turns out that the choice $b = 1$ is optimal for estimating the unit vector (direction) corresponding to the signal for any level of additive Gaussian noise before quantization as well as for a specific model of adversarial noise, while the choice $b = 2$ is optimal for estimating the direction and the norm (scale) of the signal. Moreover, Lloyd-Max quantization is shown to be an optimal quantization scheme w.r.t. ℓ_2-estimation error. Our analysis is corroborated by numerical experiments showing nearly perfect agreement with our theoretical predictions. The paper is complemented by an empirical comparison to alternative methods of signal recovery taking the non-linearity resulting from quantization into account. The results of that comparison point to a regime change depending on the noise level: in a low-noise setting, linear signal recovery falls short of more sophisticated competitors while being competitive in moderate- and high-noise settings.

1 Introduction

One of the celebrated results in compressed sensing (CS) states that it is possible to recover a high-dimensional signal $x^* \in \mathbb{R}^n$ from a small number m of Gaussian linear measurements if 1) x^* exhibits "low-dimensional structure" and 2) signal recovery is tailored to the underlying low-dimensional structure. Moreover, 2) can typically be accomplished in a computationally tractable manner, e.g., by solving a linear program. There is an enormous amount of literature on the subject spanning different areas, in particular mathematics, computer science, and engineering; we refer to [14] for an overview.

The concept of signal recovery from incomplete data in the sense of having available less measurements than would ordinarily be required has subsequently been developed further by considering settings in which the linear measurement process is subject to quantization, with the extreme case of single-bit quantization (e.g., [7, 17, 18, 22, 28, 32, 33, 44]. In general, one can think of b-bit quantization, $b \in \{1, 2, \ldots\}$. Assuming that one is free in choosing b and a corresponding scalar

quantizer given a fixed budget of bits $B = m \cdot b$ yields a trade-off between the number of measurements m and the bit depth b per measurement. An optimal balance of these two quantities minimizes a criterion of interest like the ℓ_2-error in recovering the signal. Such trade-off arises naturally in the presence of communication constraints. For example, signal acquisition and signal recovery may have to be carried out at different locations, and transmitting the acquired data is subject to a limited rate. The optimal trade-off depends on the signal, the noise mechanism and the noise level, the suitability of the parameters of the scalar quantizer relative to the signal, and the specific approach used for signal recovery. The dependency on the latter may be sidestepped by considering information-theoretical lower and upper bounds. The corresponding analysis is valuable as it would yield fundamental limits (see the survey [8] for such limits in specific settings), but it does not necessarily have immediate practical consequences unless there exists computationally tractable recovery algorithms achieving those limits. In this paper, we follow a different route by focusing on "linear" signal recovery as proposed in [33] with follow-up work in [34, 35]. Here, "linear" means that non-linearity resulting from quantization is ignored, and that the observations are treated as if they arose from a linear measurement model. Linear signal recovery may appear overly simple. In fact, it is known to be suboptimal in a noiseless setting (i.e., the only source of distortion is quantization), with an ℓ_2-error decaying with $m^{-1/2}$ compared to m^{-1} [8]. In spite of that, there is quite some justification for having a closer look at the linear approach. It turns out that ignoring non-linearity does not have a dramatic effect in a regime where the norm of the signal and the level of additive noise are comparable, in which case the ℓ_2-error of linear signal recovery can only be improved in terms of a multiplicative constant. Empirically, as is demonstrated herein, the improvements achieved by more sophisticated methods tend to be rather small. Moreover, linear signal recovery typically comes with minimum requirements in terms of computation and storage. Apart from that, linear signal recovery constitutes a natural baseline. It is thus helpful to understand the aforementioned trade-off between the number of measurements and bit depth in this simple case.

Outline and summary of contributions. In §2, we provide an overview on linear signal recovery as pioneered in [33, 34, 35] and adopt the specific formulation in [33] for estimating the "direction" $x_u^* := x^*/\|x^*\|_2$. The corresponding analysis of the trade-off between m and b when estimating x_u^* is laid out in §3. The analysis builds on ideas in [33, 35] to a considerable extent, and it is complemented by modifications/extensions to deal with the specific measurement model of b-bit quantization. Out of that analysis, we deduce explicit, easy-to-compute expressions for the relative performance of b-bit vs. b'-bit measurements ($b \neq b'$) under three different noise models (additive Gaussian noise before quantization, adversarial bin flips, and random bin flips). We are then in position to decide on the optimal choice of b given a fixed budget of bits. It turns out that the choice $b = 1$ is optimal in the noiseless case, under additive Gaussian noise as well as under adversarial bin flips. In §4, we discuss an issue that has largely been neglected in the literature, namely the estimation of the "scale" $\psi^* := \|x^*\|_2$. We show that as long as $b \geq 2$, this can be done at a fast rate by maximum likelihood estimation, separately from estimating x_u^*. Combining the results for estimating the direction and those for estimating the scale, we conclude that for the specific recovery algorithm under consideration, it does not pay off to take $b > 2$. Along the way we prove that classic Lloyd-Max quantization [29, 31] constitutes an optimal b-bit quantization scheme in the sense that it leads to a minimization of an upper bound on the ℓ_2-estimation error. Our theoretical results are corroborated by numerical experiments in §5. Regarding the relative performance of one-bit vs. two-bit measurements, the experimental results sharply agree with our theoretical predictions. One set of experiments sheds some light on the performance of linear signal recovery under study relative to alternative approaches. While the performance of the former is

noticeably inferior to more sophisticated methods in a low-noise setting, it becomes competitive as the noise level increases. Altogether, the findings of this paper point to the conclusion that noisy settings are the domain of a two-fold simple approach, consisting a basic recovery algorithm and measurements of a low bit depth (i.e., one- or two-bit measurements). More conclusions can be found in §6. Proofs and complementary derivations have been moved to the appendix.

The present paper considerably extends a previous conference publication of the authors [37]. In particular, the analysis therein is limited to sparse signals, whereas the present paper covers general low-complexity signals as quantified by the Gaussian width. In addition, we provide an account on the asymptotic sharpness of our analysis, discuss an extension to anisotropic measurements, draw additional connections to existing literature and present a more comprehensive set of numerical results.

Related work. There is a plethora of papers discussing various aspects of compressed sensing with quantization. We refer to [8] for an excellent overview on the problem including basic performance limits, different approaches to quantization and signal recovery, and the associated references. On the other hand, comparatively little seems to be known about the trade-off between the number of measurements and bit depth as it is in the focus of the present paper. An important reference in this regard is [25] where this very trade-off is studied for sparse signals. The analysis in [25] concerns "oracle-assisted" least squares (i.e., least squares with knowledge of the set of non-zero entries of the signal) which is of theoretical interest, but not a practical approach to signal recovery. The authors point out the role of the signal-to-noise ratio (SNR) for the optimal trade-off that leads to the distinction of two basic regimes: the so-called measurement compression regime with high SNR, small m and large b as opposed to the quantization compression regime with low SNR, large m and small b. Some of the results in our paper can be related to this finding.

We here focus on linear signal recovery. An alternative that is more principled as it uses knowledge about the quantizer is "consistent reconstruction". This approach has been studied in a series of recent papers by Jacques and collaborators [21, 1, 20]. At the moment, this line of research only addresses the case without noise. The trade-off between m and b appears in [21], where a specific version of Iterative Hard Thresholding [5] tailored to quantized measurements is considered. It is shown via experimental results (which are partially reproduced in §5) that unlike the main result of the present paper, increasing b beyond one or two yields improvements, which underlines that the trade-off can be rather different depending on the recovery algorithm used.

After submitting the conference paper [37], we became aware of the work [38] in which Lloyd-Max quantization is found to be optimal for linear signal recovery as in the present paper. The derivation in [38] is not fully rigorous though as it is only shown that Lloyd-Max quantization yields a stationary point, whereas herein, global optimality is established.

Notation. For the convenience of the reader, we here gather notation used throughout the paper. For a positive integer d, we use the shortcut $[d] = \{1, \dots, d\}$. $I(P)$ denotes the indicator function of expression P with $I(P) = 1$ if P is true and 0 otherwise. For a matrix $A = (a_{ij})_{1\le i\le m, 1\le j\le n}$, a_i denotes the i-th row and A_j the j-th column, $i \in [m]$, $j \in [n]$. We use $\|A\|$ for the spectral norm of A. The identity matrix of dimension d is denoted by I_d. For a vector x, we write x_J for the sub-vector corresponding to an index set J. For $x, x' \in \mathbb{R}^n$, $x \odot x' = (x_j \cdot x'_j)_{j=1}^n$. A class of signals is denoted by $\mathcal{K} \subset \mathbb{R}^n$ and we let $\mathcal{C} = \mathcal{K} \cap B_2^n$, where for $q \ge 1$, B_q^n denotes the unit ℓ_q-ball in $\mathbb{R}^n$. We further write $B_0(r; n) = \{x \in \mathbb{R}^n : \|x\|_0 \le r\}$, where $\|x\|_0 = \sum_{j=1}^n I(x_j \neq 0)$. The unit sphere of $\mathbb{R}^n$ is denoted by $\mathbb{S}^{n-1}$. For a set $K \subseteq \mathbb{R}^n$, we write $|K|$ for its cardinality, $\mathrm{conv}(K)$ for its convex hull, and for $a \in \mathbb{R}$, we let $aK = \{a \cdot x, x \in K\}$. The letter g refers to a Gaussian random variable or a canonical Gaussian random vector, i.e., $g \sim N(0, 1)$ or $g \sim N(0, I_n)$. The probability density function (pdf) and the cumulative density function (cdf) of the standard Gaussian distribution are

denoted by ϕ and Φ, respectively. In addition to the usual Landau notation, we occasionally make use of the stochastic order symbol $O_{\mathbf{P}}$: a sequence of random variables (ε_n) satisfies $\varepsilon_n = O_{\mathbf{P}}(g(n))$ if for all $\delta > 0$ there is a finite c such that $\mathbf{P}(|\varepsilon_n/g(n)| > c) \leq \delta$ for all n.

2 Linear signal recovery based on quantized linear measurements

In this section, we first fix the problem setup and then introduce the approach that will be studied in depth in subsequent sections.

Measurement model. Let $x^* \in \mathcal{K} \subseteq \mathbb{R}^n$ be the signal to be recovered. We think of $\mathcal{K}$ as a set describing a class of signals having a certain low-dimensional structure, e.g., $\mathcal{K} = B_0(s;n) := \{x \in \mathbb{R}^n : \|x\|_0 \leq s\}$, the set of s-sparse signals. More examples are given in §3.3. The set $\mathcal{K}$ is assumed to be known. Let $A = (a_{ij})_{1\leq i\leq m, 1\leq j\leq n}$ be a random matrix with i.i.d. $N(0,1)$ entries whose rows and columns are denoted by $\{a_i\}_{i=1}^m$ and $\{A_j\}_{j=1}^n$, respectively. The observations $y = (y_i)_{i=1}^m$ arise from the model

$$y_i = Q(\langle a_i, x^* \rangle + \sigma\varepsilon_i), \;\; i = 1, \ldots, m, \tag{1}$$

where $\varepsilon = (\varepsilon_i)_{i=1}^m$ has i.i.d. $N(0,1)$ entries and $\sigma \geq 0$ is referred to as "noise level"; when $\sigma = 0$, we speak of a noiseless setting. The map $Q : \mathbb{R} \to \mathcal{M}^{\pm}$ is called quantization map or quantizer, which is piecewise constant, monotonically increasing, and odd. It partitions the real axis into 2^b bins where $b \in \{1, 2, \ldots\}$ is the bit depth per measurement. Because of symmetry, it suffices to define a partitioning of $\mathbb{R}_+$ into $K = 2^{b-1}$ bins $\{\mathcal{R}_k\}_{k=1}^K$ resulting from distinct thresholds $\mathbf{t} = (t_1, \ldots, t_{K-1})^\top$ (in increasing order) and $t_0 = 0$, $t_K = +\infty$ such that $\mathcal{R}_1 = [t_0, t_1), \ldots, \mathcal{R}_K = [t_{K-1}, t_K)$. Each bin is assigned a distinct representative from the codebook $\mathcal{M}^{\pm} = -\mathcal{M} \cup \mathcal{M}$ with $\mathcal{M} = \{\mu_1, \ldots, \mu_K\}$. Accordingly, Q is defined as

$$z \mapsto Q(z) = \operatorname{sign}(z) \sum_{k=1}^{K} \mu_k I(|z| \in \mathcal{R}_k). \tag{2}$$

For convenience, Figure 1 visualizes this definition. Since noise (if any) is added before Q is applied, we speak of "additive noise before quantization". Other noise mechanism acting after quantization are possible, too; see §3.5 for specific examples.

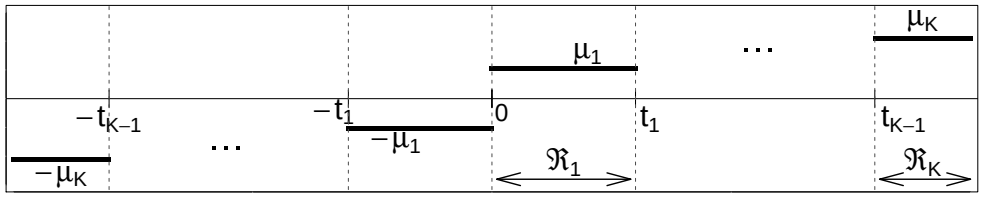

Figure 1: Schematic representation of the quantization map and the associated quantities. Note that it suffices to define Q via a partitioning of $\mathbb{R}_+$, which then extends to a partitioning of $\mathbb{R}$ (as shown here) by symmetry. The map Q is a step function with levels $-\mu_K, \ldots, -\mu_1, \mu_1, \ldots, \mu_K$ and breakpoints $-t_{K-1}, \ldots, -t_1, t_0, t_1, \ldots, t_{K-1}$.

Linear Signal Recovery. Linear signal recovery is carried out by means of an estimator $\widehat{\theta}(A, y) = \widehat{\theta}(\mathcal{L}_A y)$ for some linear map $\mathcal{L}_A : \mathbb{R}^m \to \mathbb{R}^n$ depending on A and $\widehat{\theta}$ only depending on $\mathcal{K}$ but not directly on Q. It is perhaps surprising that even when restricting oneself to this class, it is possible to construct consistent estimators of the "direction" $x_u^* = x^*/\|x^*\|_2$. This was first established by Brillinger [9] in the traditional setting of asymptotic statistics with fixed parameter set and m tending to infinity. It turns out that Gaussianity of A plays a crucial role here. Brillinger's result

has been generalized recently to modern high-dimensional settings in various ways [33, 34, 35, 38]. We follow this line of research in the present paper. Linear signal recovery is limited to estimating x_u^* as estimating the "scale" $\psi^* = \|x^*\|_2$ entails using knowledge of Q. Indeed, expanding model (1), we obtain

$$\begin{aligned} y_i &= \operatorname{sign}(\langle a_i, x^* \rangle + \sigma\varepsilon_i) \textstyle\sum_{k=1}^K \mu_k I(\,|(\langle a_i, x^* \rangle + \sigma\varepsilon_i)| \in \mathcal{R}_k) \\ &= \operatorname{sign}(\langle a_i, x_u^* \rangle + \tau\varepsilon_i) \textstyle\sum_{k=1}^K \mu_k I(\,|(\langle a_i, x_u^* \rangle + \tau\varepsilon_i)| \in \mathcal{R}_k/\psi^*),\ i \in [m], \end{aligned} \tag{3}$$

where $\tau = \sigma/\psi^*$. We conclude that model (1) can always be transformed into one with $\psi^* = 1$ by re-scaling the thresholds $\mathbf{t}$ accordingly. Hence, in order to be able to estimate ψ^*, the thresholds $\mathbf{t}$ have to be taken into account, while linear signal recovery discards this information. Moreover, a second consequence is that when studying linear signal recovery with regard to the estimation of x_u^* in the sequel, we may assume w.l.o.g. that $\psi^* = 1$. Linear estimation of the direction can trivially be combined with separate non-linear estimation of the scale, an approach whose discussion is postponed to §4.

Marginal Regression. The choice $\mathcal{L}_A = m^{-1}A^\top$ gives rise to what we will refer to as the "canonical linear estimator":

$$\widehat{x} \in \operatorname*{argmin}_{x \in \mathcal{C}} -\langle \eta, x \rangle\,, \quad \text{where } \eta := \frac{A^\top y}{m}, \quad \text{and } \ \mathcal{C} := \mathcal{K} \cap B_2^n. \tag{4}$$

This estimator is proposed in [33] for $\mathcal{K} = \sqrt{s}B_1^n$ and $b = 1$. In the conference paper preliminary to the present work [37], $\mathcal{K} = B_0(s; n)$ and general b are considered. The heading "Marginal Regression" is rooted in the fact that η is essentially proportional to the vector of univariate (or marginal) regression coefficients $\widehat{\beta} = (\widehat{\beta}_j)_{j=1}^n$ corresponding to n separate linear regressions of y onto A_j, $j \in [n]$, which are given by

$$\widehat{\beta}_j = \frac{A_j^\top y}{\|A_j\|_2^2} = \frac{A_j^\top y}{m + O_{\mathbf{P}}(\sqrt{m})}$$

which implies that $\eta_j = (1 + O_{\mathbf{P}}(m^{-1/2}))\widehat{\beta}_j$, $j \in [n]$.

A formulation rather similar to (4) is studied in [35].

$$\widehat{x}_\lambda \in \lambda^{-1}\Pi_{\lambda\mathcal{K}}(\eta) = \lambda^{-1} \operatorname*{argmin}_{x \in \lambda\mathcal{K}} \frac{1}{2}\|\eta - x\|_2^2 = \lambda^{-1} \operatorname*{argmin}_{x \in \lambda\mathcal{K}} \frac{1}{2}\|x\|_2^2 - \langle \eta, x \rangle \tag{5}$$

where $\lambda > 0$ is defined by the relation $\mathbf{E}[A^\top y/m] = \lambda x_u^*$ (cf. (7) below) and $\Pi_{\lambda\mathcal{K}}$ denotes the Euclidean projection on $\lambda\mathcal{K}$. As derived in Appendix H, if $\mathcal{K}$ is a cone (i.e., $\alpha\mathcal{K} = \mathcal{K}$ for all $\alpha > 0$), then $\widehat{x}_\lambda$ is identical to $\widehat{x}$ in (4) up to a constant of proportionality. This applies to all except for one of the examples considered for $\mathcal{K}$ herein. In general, $\widehat{x}/\|\widehat{x}\|_2$ and $\widehat{x}_\lambda/\|\widehat{x}_\lambda\|_2$ are different though. A notable disadvantage of (5) compared to (4) is that it requires knowledge of λ which depends on the (typically unknown) noise level σ. For this reason, we concentrate on (4) in the following.

Least Squares. The choice $\mathcal{L}_A = (A^\top A)^{-1}A^\top$ would yield the least squares estimator. Since the inverse does not exist once $m < n$, one needs to take advantage of low-dimensional structure. This yields the constrained least squares estimator

$$\widehat{x}_{\mathrm{LS}} \in \lambda^{-1} \operatorname*{argmin}_{x \in \lambda\mathcal{K}} \frac{1}{2} x^\top \frac{A^\top A}{m} x - \langle \eta, x \rangle\,. \tag{6}$$

In [34] and [38], this approach is studied under the names "Generalized Lasso" respectively "$\mathcal{K}$-Lasso" in allusion to the popular choice for the set $\mathcal{K}$ as an ℓ_1-ball [39]. Observe that (5) differs from (6) only in that the matrix $A^\top A/m$ is replaced by its expectation, the identity matrix. While this may appear as minor, an important consequence is that (6) can achieve *exact* recovery of x^* as $b \to \infty$ and $\sigma = 0$ unlike (5). Outside the high signal-to-noise regime, however, both formulations indeed perform similarly as can be concluded from the analysis of (5) in [35] and the analysis of (6) in [34]. The latter yields a bound that is essentially of the form $\|\widehat{x}_{\mathrm{LS}} - x_u^*\|_2 \leq c^{-1}\|\widehat{x}_\lambda - x_u^*\|_2$ resulting from a lower bound of the form $\frac{1}{m}\|Ax\|_2^2 \geq c\|x\|_2^2$, $0 < c < 1$, for all x in the so-called tangent cone of $\lambda\mathcal{K}$ at λx^*. There are two more aspects that are relevant to a comparison.

- From the viewpoint of computation, the optimization problems (4) and (5) tend to be easier than the one in (6) as the objectives in (4),(5) are linear respectively coordinate-wise separable.
- More generally, one can consider the case of i.i.d. anisotropic measurements, i.e., $a_i \sim N(0, \Sigma)$, $i \in [m]$. If Σ is known, then (4) and (5) remain applicable with $\mathcal{L}_A = m^{-1}\Sigma^{-1}A^\top$, cf. §3.4 below. The approach (6) has its merits in the situation that Σ is not known. In general, estimating Σ resp. its inverse is statistically more difficult and computationally more demanding than estimating x^*, hence the use of (4) or (5) combined with plug-in estimation of Σ^{-1} is not a suitable option.

3 Analysis

The following section is dedicated to the analysis of the ℓ_2-error of the canonical linear estimator (4) under the b-bit quantization model as defined by (1) and (2). Our main result is an asymptotic bound for $n \to \infty$ that allows for a precise (asymptotically sharp) characterization of the dependence on the bit depth b, the thresholds $\mathbf{t}$ and the representatives $\boldsymbol{\mu} = (\mu_1, \dots, \mu_K)^\top$ parameterizing the quantization map. Given this result, we are in position to address the trade-off between m and b. Along the way, we show that Lloyd-Max quantization [29, 31] constitutes an optimal quantization scheme in the sense that it leads to a minimization of the error bound w.r.t. $\mathbf{t}$ and $\boldsymbol{\mu}$. Finally, an extension to two other natural noise models is discussed.

At a technical level, the main ingredients of our analysis appear in related literature, in particular in [33]. Certain adjustments are necessary though to deal effectively with the specific measurement model herein, and to end up with a result suitable for the purpose of studying the trade-off between m and b.

3.1 Preparations

Our main result depends on three quantities and one condition which are given below. Throughout this section, we assume w.l.o.g. that $\psi^* = 1$ so that $x^* = x_u^*$ as this can always be achieved by re-scaling $\mathbf{t}$ and σ, cf. (3).

(Q1) The first quantity has initially been introduced in [33].

$$\lambda = \lambda_b = \lambda_{b,\sigma} = \lambda_{b,\sigma}(\mathbf{t}, \boldsymbol{\mu}) := \mathbf{E}[\theta(g)g], \quad g \sim N(0,1), \tag{7}$$

where the map θ is defined by the relation

$$\mathbf{E}[y_1|a_1] = \theta(\langle a_1, x_u^*\rangle). \tag{8}$$

At a high level, λ quantifies the distortion from linearity caused by quantization. It can be shown that (Appendix A) that $\mathbf{E}[A^\top y/m] = \lambda x_u^*$, hence λ also equals the constant of proportionality up

to which x_u^* can be recovered by linear estimation. The quantity λ is positive, increases with b and approaches one as $b \to \infty$. A precise expression for $\lambda_{b,\sigma}(\mathbf{t}, \boldsymbol{\mu})$ is the content of Lemma 1 below.

(Q2) The second quantity is given by

$$\Psi = \Psi_{b,\sigma} = \Psi_{b,\sigma}(\mathbf{t}, \boldsymbol{\mu}) := \sqrt{\mathbf{E}[y_1^2]}, \tag{9}$$

which is simply the (marginal) standard deviation of the $\{y_i\}_{i=1}^m$. The error bound of Theorem 1 below is proportional to Ψ. We refer to Lemma 1 for a more specific expression for $\Psi_{b,\sigma}(\mathbf{t}, \boldsymbol{\mu})$.

(Q3) The third quantity is a measure of complexity of the class of signals $\mathcal{K}$ under consideration, the so-called Gaussian width of the tangent cone of $\mathcal{C} = \mathcal{K} \cap B_2^n$ at x_u^*, a notion which can be considered as standard in the context of high-dimensional linear inverse problems [10, 13, 33, 34, 35]. We set

$$\Delta(\mathcal{C}; x_u^*) := \{\alpha(x - x_u^*),\ \alpha \geq 0,\ x \in \mathcal{C}\}, \qquad \overline{\Delta}(\mathcal{C}; x_u^*) := \Delta(\mathcal{C}; x_u^*) \cap \mathbb{S}^{n-1}. \tag{10}$$

We suppress dependence on x_u^*, i.e., we use $\Delta(\mathcal{C}) = \Delta(\mathcal{C}; x_u^*)$ and $\overline{\Delta}(\mathcal{C}) = \overline{\Delta}(\mathcal{C}; x_u^*)$ for the tangent cone of $\mathcal{C}$ at x_u^* and its spherical part, respectively. The latter enters Theorem 1 below via its Gaussian width. For $K \subset \mathbb{R}^n$ compact and $g \sim N(0, I_n)$, the Gaussian width of K is defined by

$$w(K) := \mathbf{E}\left[\sup_{v \in K} |\langle g, v \rangle|\right] \tag{11}$$

(C) Finally, we require the following condition. At a technical level, rather than being fixed, we think of x_u^* as a sequence $\{x_u^{*,n}\}_{n \in \mathbb{N}}$ whose elements are contained in spheres of growing dimension, satisfying $\|x_u^{*,n}\|_\infty \to 0$ as $n \to \infty$. This condition is easily met for typical signal classes of interest under natural sampling models for x_u^* as is elaborated in the discussion after Theorem 1 below.

3.2 Main result

We now state our main result along with a brief general discussion. Further implications are subsequently discussed in separate subsections.

Theorem 1. *Consider the measurement model given by* (1),(2) *and consider the canonical linear estimator* $\widehat{x}$ *given by* (4). *Under condition* **(C)**, *as* $n \to \infty$, *we have*

$$\|\widehat{x} - x_u^*\|_2 \leq 24\, \frac{\Psi_{b,\sigma}(\mathbf{t}, \boldsymbol{\mu})}{\lambda_{b,\sigma}(\mathbf{t}, \boldsymbol{\mu})}\, \frac{w(\overline{\Delta}(\mathcal{C}))}{\sqrt{m}}. \tag{12}$$

with probability at least $1 - 4(\exp(-w(\overline{\Delta}(\mathcal{C}))/2) + \exp(-cm)) = 1 - o(1) - 4\exp(-cm)$ *as* $n \to \infty$, *for some constant* $c > 0$.

It turns out that the bound (12) does not leave much room for further improvement in general. The term $w(\overline{\Delta}(\mathcal{C}))/\sqrt{m}$ is identified as the rate of estimation. It follows from existing literature [10, 13] that in the presence of additive noise, one has a corresponding lower bound in terms of the so-called Sudakov minoration of $w(\overline{\Delta}(\mathcal{C}))$ which typically yields matching upper and lower bounds modulo constant factors. For the examples of $\mathcal{K}$ resp. $\mathcal{C}$ given in §3.3 below, the upper bounds on $w(\overline{\Delta}(\mathcal{C}))$ yield a match to known minimax lower bounds, e.g., [11, 12, 30].

We point out, however, that the rate in (12) can be suboptimal in the noiseless case. For $\mathcal{K} = B_0(s; n)$, it is shown in [22] that there exists a (computationally intractable) recovery algorithm that achieves an error decay of the order m^{-1} compared to $m^{-1/2}$ in (12).

In the sequel, a lot of attention will be paid to the leading constant $\Psi_{b,\sigma}(\mathbf{t}, \boldsymbol{\mu})/\lambda_{b,\sigma}(\mathbf{t}, \boldsymbol{\mu})$. The dependency on this quantity is asymptotically sharp as it is already encountered in the traditional

asymptotic setup in which $\mathcal{K} = \mathbb{R}^n$ with n being of a smaller order of magnitude than m. In fact, one can show (cf. Appendix D) that under a double asymptotic framework in which $m, n \to \infty$, $n/m \to 0$ and **(C)** holds, for any $j \in [n]$

$$\sqrt{m}\left(\widehat{x}_j - x^*_{u,j}\right) \xrightarrow{\mathcal{D}} N\left(0, \frac{\Psi^2_{b,\sigma}(\mathbf{t}, \boldsymbol{\mu})}{\lambda^2_{b,\sigma}(\mathbf{t}, \boldsymbol{\mu})}\right), \tag{13}$$

where $\xrightarrow{\mathcal{D}}$ denotes convergence in distribution. In other words, the estimation error for any single coordinate is proportional to the leading constant of our bound.

The numerical constant "24" in the bound (12) does not appear to be optimal. As it comes to the key point of the paper, namely the ratio of estimation errors for different choices of the bit depth b, this is not an issue as all terms not depending on b cancel out.

3.3 Classes of signals

Theorem 1 can be specialized to popular signal classes by bounding the associated Gaussian widths. We here provide several examples including a short discussion of computational aspects. We also discuss condition (**C**) in light of those examples. Derivations have been relegated to Appendix J.

1) *Sparsity.* $\mathcal{K} = B_0(s; n)$

Following the argument in the proof of Lemma 2.3 in [33], one can show that

$$w(\overline{\Delta}(\mathcal{C})) \leq w(B_0(2s; n) \cap \mathbb{S}^{n-1}) \leq \sqrt{2s} + \sqrt{4s \log\left(\frac{en}{2s}\right)} + \sqrt{\pi/2} \leq 3.5\sqrt{2s \log\left(\frac{en}{2s}\right)},$$

i.e., we recover the usual rate $\left(\frac{s \log(n/s)}{m}\right)^{1/2}$ in Theorem 1.

2) *Fused Sparsity.* Let $D : \mathbb{R}^n \to \mathbb{R}^{n-1}$ denote the first-order difference operator and set $\mathcal{K} = \mathrm{PC}(s; n) := \{x \in \mathbb{R}^n : Dx \in B_0(s; n-1)\}$; this is the set of all signals in $\mathbb{R}^n$ that are piecewise constant with s breakpoints. One can show that $w(\overline{\Delta}(\mathcal{C}))$ satisfies the same upper bound as in 1).

3) *Group Sparsity.* Let $\mathcal{G} = \{G_\ell\}_{\ell=1}^L$ be a partition of $[n]$ into L groups. Define $\|\cdot\|_{0,\mathcal{G}}$ by $|\{\ell \in [L] : x_{G_\ell} \neq 0\}|$ and for $1 \leq r \leq L$, let $B_{0,\mathcal{G}}(r) = \{x \in \mathbb{R}^n : \|x\|_{0,\mathcal{G}} \leq r\}$. Consider $\mathcal{K} = B_{0,\mathcal{G}}(s)$, in which case we say that x^* is s-group sparse (w.r.t. the partition $\mathcal{G}$). Then, $w(\overline{\Delta}(\mathcal{C})) \leq 2\sqrt{2s \max_{\ell \in [L]} |G_\ell|} + \sqrt{4s \log(eL/2s)}$.

4) *Low-rank matrices.* Our framework can accommodate a set of matrices $\mathcal{K} \subseteq \mathbb{R}^{n_1 \times n_2}$ by identifying $\mathbb{R}^{n_1 \times n_2}$ with $\mathbb{R}^{n_1 n_2}$. Consider $\mathcal{K} = \{X \in \mathbb{R}^{n_1 \times n_2} : \mathrm{rank}(X) \leq s\}$ and accordingly $\mathcal{C} = \mathcal{K} \cap \{X \in \mathbb{R}^{n_1 \times n_2} : \|X\|_F \leq 1\}$, with $\|\cdot\|_F$ as the Frobenius norm. Then, $w(\overline{\Delta}(\mathcal{C})) \leq \sqrt{2sn_1} + \sqrt{2sn_2}$.

5) *ℓ_1-ball constraint.* $\mathcal{K} = \mathcal{C} = r^* B_1^n \cap B_2^n$, $r^* = \|x^*_u\|_1$. In addition, suppose that $x^* \in B_0(s; n)$. Then $w(\overline{\Delta}(\mathcal{C})) \leq 2\sqrt{2} w(B_0(2s; n) \cap \mathbb{S}^{n-1})$, and we may resort to the bound in 1).

In the same way as 5) arises as the convex counterpart to 1), one can consider a total variation constraint in place of 2), an $\ell_1 - \ell_2$ ball ("group lasso" [43]) constraint in place of 3), and a Schatten-one norm ball constraint in place 4). Under a sparsity assumption for x^*, the Gaussian widths of the convex formulations equal – up to numerical constants – those of the corresponding non-convex formulations; for the sake of brevity and since this is known in the literature, we omit explicit statements/derivations here.

Computation. The constraint sets in examples 1) to 4) are non-convex. Nevertheless, due to the simplicity of the objective in Eq. (4) and the specific structure of the constraint sets, all of

the resulting optimization problems are computationally tractable; 1) and 3) even have closed form solutions, 4) can be reduced to a singular value decomposition, and 2) can be solved in $O(n^2 s)$ flops by dynamic programming [3]. A formal derivation for 1) is contained in Appendix I representative for 1) – 3). For 2), the convex counterpart has a much better computational complexity which scales only linearly in n [24], while achieving comparable statistical performance.

Discussion of **(C)**. For examples 1) – 4) above, Condition **(C)** can be shown to be satisfied with probability tending to one as $n \to \infty$ when sampling x_u^* uniformly at random according to natural generating mechanisms.

1) For a given sparsity level s, pick the support of x^* at random and sample the non-zero entries i.i.d. from a distribution with finite fourth moment, and normalize to unit 2-norm. It follows from Markov's inequality that **(C)** is satisfied with probability tending to one as $n, s \to \infty$.

2) Pick a random partition $\{1, \dots, i_{k_1}\}, \{i_{k_1+1}, \dots, i_{k_2}\}, \dots, \{i_{k_{s-1}+1}, \dots, n\}$ of $\{1, \dots, n\}$ such that $n_k/n \to c_k$, $\sum_{k=1}^{s} c_k = 1$ as $n \to \infty$, where n_k denotes the number of elements in the k-th element of the partition, $k \in [s]$. For each of those, sample the corresponding entries at random from a distribution with finite first moment, scale them by the square root of the respective block size and then normalize to unit 2-norm.

3) Suppose that as $n \to \infty$, the sizes of the non-zero blocks are of the same order. Sampling the non-zero entries as in example 1), condition **(C)** is fulfilled with probability tending to one as the block sizes of the non-zero blocks or the number of non-zero blocks go to infinity.

4) Draw a random $n_1 \times n_2$ matrix with i.i.d. $N(0,1)$-entries, $n_1, n_2 > r$, and compute its SVD. Keep the top r left and right singular vectors and replace the corresponding singular values by an arbitrary element of $\mathbb{S}^{r-1}$.

Clearly, there may exist different or more general sampling schemes for **(C)** to be satisfied.

3.4 Extension to anisotropic measurements

We extend Theorem 1 to the case of anisotropic Gaussian measurements. More precisely, we now suppose that the $\{a_i\}_{i=1}^m$ are i.i.d. from a $N(0, \Sigma)$-distribution where Σ is invertible and assumed to be known. In the anisotropic case, the linear estimator (4) is replaced by

$$\widehat{x}_\Sigma \in \operatorname*{argmin}_{x \in \mathcal{C}} - \left\langle x, \Sigma^{-1} \eta \right\rangle, \tag{14}$$

with $\eta = A^\top y/m$ as in (4). We then have the following counterpart to Theorem 1.

Theorem 2. *Consider the anisotropic measurement model as above, let $\kappa(\Sigma^{1/2})$ denote the condition number of $\Sigma^{1/2}$, and let $\widehat{x}_\Sigma$ be as in* (14)*. Under condition* **(C)***, as $n \to \infty$, it holds that*

$$\|\widehat{x}_\Sigma - x_u^*\|_2 \leq 24 \kappa(\Sigma^{1/2}) \, \frac{\Psi_{b,\sigma}(\mathbf{t}, \boldsymbol{\mu})}{\lambda_{b,\sigma}(\mathbf{t}, \boldsymbol{\mu})} \, \frac{w(\overline{\Delta}(\mathcal{C}))}{\sqrt{m}}, \tag{15}$$

with probability at least $1 - 4(\exp(-\kappa(\Sigma^{1/2}) w(\overline{\Delta}(\mathcal{C}))/2) + \exp(-cm)) = 1 - o(1) - 4\exp(-cm)$ as $n \to \infty$, for some constant $c > 0$.

The bound for the anisotropic case thus only involves the additional factor $\kappa(\Sigma^{1/2})$. Setting $\Sigma = I_n$, we recover Theorem 1. Regarding the trade-off between m and b to be studied in the next sections, the extra factor does not have any influence as it does not depend on b.

3.5 Implications for the optimal trade-off between m and b

We now study in detail the implications of Theorem 1 for the central question of this paper. Suppose we have a fixed budget of B bits available and are free to choose the number of measurements m and the number of bits per measurement b subject to $B = m \cdot b$ such that the ℓ_2-error of $\|\widehat{x} - x_u^*\|_2$ of is as small as possible. What is the optimal choice of (m, b)? At this point, we still confine ourselves to the direction x_u^*. In §4 below, we provide an answer for x^* in place of x_u^* by linking the findings of the present section to the results on scale estimation.

In virtue of Theorem 1 and the comments that follow, the asymptotic ℓ_2-error as $n \to \infty$ depends on b only via $\Omega_{b,\sigma}(\mathbf{t}, \boldsymbol{\mu}) := \Psi_{b,\sigma}(\mathbf{t}, \boldsymbol{\mu})/\lambda_{b,\sigma}(\mathbf{t}, \boldsymbol{\mu})$. In a first step, we provide more specific expressions for $\Psi_{b,\sigma}(\mathbf{t}, \boldsymbol{\mu})$ and $\lambda_{b,\sigma}(\mathbf{t}, \boldsymbol{\mu})$ that will prove useful in subsequent analysis. Below, $\odot$ denotes the entry-wise (Hadamard) multiplication of vectors.

Lemma 1. *Let $\mathcal{R}_k = \mathcal{R}_k(\mathbf{t})$, $k \in [K]$, be as in (2). We have $\lambda_{b,\sigma}(\mathbf{t}, \boldsymbol{\mu}) = \langle \boldsymbol{\alpha}(\mathbf{t}), \boldsymbol{E}(\mathbf{t}) \odot \boldsymbol{\mu} \rangle / (1+\sigma^2)$ and $\Psi_{b,\sigma}(\mathbf{t}, \boldsymbol{\mu}) = \sqrt{\langle \boldsymbol{\alpha}(\mathbf{t}), \boldsymbol{\mu} \odot \boldsymbol{\mu} \rangle}$, where*

$$\begin{aligned}
\boldsymbol{\alpha}(\mathbf{t}) &= (\alpha_1(\mathbf{t}), \dots, \alpha_K(\mathbf{t}))^\top, \quad \alpha_k(\mathbf{t}) = \mathbf{P}\left\{|\widetilde{g}| \in \mathcal{R}_k(\mathbf{t})\right\}, \ \widetilde{g} \sim N(0, 1+\sigma^2), \ \ k \in [K], \\
\boldsymbol{E}(\mathbf{t}) &= (E_1(\mathbf{t}), \dots, E_K(\mathbf{t}))^\top, \ E_k(\mathbf{t}) = \mathbf{E}[\widetilde{g} | \widetilde{g} \in \mathcal{R}_k(\mathbf{t})], \quad \widetilde{g} \sim N(0, 1+\sigma^2), \ \ k \in [K].
\end{aligned}$$

Optimal choice of $\mathbf{t}$ and $\boldsymbol{\mu}$. In order to eliminate the dependence on $\mathbf{t}$ and $\boldsymbol{\mu}$, we minimize $\Omega_{b,\sigma}$ w.r.t. these two quantities. In this manner, we also obtain an optimal parameterization of the quantization map Q yielding minimum ℓ_2-estimation error. It turns out that the solution coincides with that of the classical Lloyd-Max quantization problem [29, 31] stated below. Let h be a random variable with finite variance and consider the optimization problem

$$\min_{\mathbf{t}, \boldsymbol{\mu}} \mathbf{E}[\{h - Q(h; \mathbf{t}, \boldsymbol{\mu})\}^2] = \min_{\mathbf{t}, \boldsymbol{\mu}} \mathbf{E}[\{h - \operatorname{sign}(h) \textstyle\sum_{k=1}^K \mu_k I(|h| \in \mathcal{R}_k(\mathbf{t}))\}^2]. \tag{16}$$

Problem (16) can be solved by an iterative scheme known as the Lloyd-Max algorithm (§3.2.3 in [16]) that alternates between optimization of $\mathbf{t}$ for fixed $\boldsymbol{\mu}$ and vice versa. For h from a log-concave distribution (e.g., Gaussian) that scheme can be shown to deliver the global optimum [23].

Theorem 3. *Consider the minimization problem $\min_{\mathbf{t}, \boldsymbol{\mu}} \Omega_{b,\sigma}(\mathbf{t}, \boldsymbol{\mu})$ with $\Omega_{b,\sigma}(\mathbf{t}, \boldsymbol{\mu}) = \frac{\Psi_{b,\sigma}(\mathbf{t}, \boldsymbol{\mu})}{\lambda_{b,\sigma}(\mathbf{t}, \boldsymbol{\mu})}$. Its minimizer $(\mathbf{t}^*, \boldsymbol{\mu}^*)$ equals that of the Lloyd-Max problem (16) for $h \sim N(0, 1+\sigma^2)$. Moreover,*

$$\Omega_{b,\sigma}(\mathbf{t}^*, \boldsymbol{\mu}^*) = \sqrt{(\sigma^2+1)/\lambda_{b,0}(\boldsymbol{t}_0^*, \boldsymbol{\mu}_0^*)}, \tag{17}$$

where $(\boldsymbol{t}_0^*, \boldsymbol{\mu}_0^*)$ denotes the optimal choice of $(\mathbf{t}, \boldsymbol{\mu})$ in the sense of (16) when $h \sim N(0, 1)$ (i.e., $\sigma = 0$).

Regarding the choice of $(\mathbf{t}, \boldsymbol{\mu})$ the result of Theorem 3 may not come as a surprise as the entries of y are i.i.d. $N(0, 1+\sigma^2)$. It is less immediate though that this specific choice can also be motivated as the one leading to the minimization of the error bound (12).

The second part of Theorem 3 implies that the ratio $\Omega_{b,\sigma}(\mathbf{t}^*, \boldsymbol{\mu}^*)/\Omega_{b',\sigma}(\mathbf{t}^*, \boldsymbol{\mu}^*)$ does not depend on σ. Combining Theorem 1, Lemma 1 and Theorem 3, we are eventually in position to determine the optimal trade-off between m and b. Theorem 1 yields that the ℓ_2-error decays with $m^{-1/2}$. Therefore, for b' to improve over b with $b < b'$ at the level of bits, it is required that $\Omega_b/\Omega_{b'} > \sqrt{b'/b}$. In fact, when using b bits per measurement we may multiply the number of measurements by a factor of b'/b so that the bit budgets are balanced, i.e., $m_b \cdot b = m_{b'} \cdot b'$, where m_b and $m_{b'}$ denote the number of b-bit and b'-bit measurements, respectively. Using the Lloyd-Max algorithm to determine

	$b=1,\ b'=2$	$b=2,\ b'=3$	$b=3,\ b'=4$
$\Omega_b/\Omega_{b'}$:	**1.178**	**1.046**	**1.013**
required for $b' \gg b$:	$\sqrt{2} \approx$ **1.414**	$\sqrt{3/2} \approx$ **1.225**	$\sqrt{4/3} \approx$ **1.155**

Table 1: Here $b' \gg b$ means that b' is a superior choice relative to b at the bit level.

$(\mathbf{t}^*, \boldsymbol{\mu}^*)$ and invoking (17) as well as Lemma 1, it is straightforward to evaluate the ratios $\Omega_b/\Omega_{b'}$ numerically. In Table 1, we provide the results for selected pairs of b and b'.

From these figures, we see that increasing b reduces the error as the number of measurements are fixed. However, the reduction is not substantial enough to yield an improvement when thinking in terms of a budget of bits instead of measurements. Reducing b from two to one increases the error by a factor of 1.178 which is below $\sqrt{2}$, the factor required for $b=1$ to be inferior compared to $b=2$. The reduction factor for increasing b becomes even smaller for the transitions from two to three and three to four bits, and quickly approaches 1. The overall conclusion is that for estimating x_u^* the optimal trade-off is achieved by one-bit quantization – instead of increasing the bit depth b, one should rather increase the number of measurements. This conclusion is valid regardless of the noise level as a consequence of Theorem 3. Even more, the figures in Table 1 assume optimal quantization for $b \geq 2$, and in turn knowledge of σ, which may not be fulfilled in practice.

Beyond additive noise. Additive Gaussian noise is perhaps the most studied form of perturbation, but one can of course think of numerous other mechanisms whose effect can be analyzed along the path used for additive noise as long as it is feasible to obtain the corresponding expressions for λ and Ψ. We here do so for the following mechanisms acting *after* quantization (2).
(I) *Random bin flip.* For $i \in [m]$: with probability $1-p$, y_i remains unchanged. With probability p, y_i is changed to an element from $\mathcal{M}^{\pm} \setminus \{y_i\}$ uniformly at random.
(II) *Adversarial bin flip.* For $i \in [m]$: write $y_i = q\mu_k$ for $q \in \{-1, 1\}$ and $\mu_k \in \mathcal{M}$. With probability $1-p$, y_i remains unchanged. With probability p, y_i is changed to $-q\mu_K$.

Note that for $b=1$, (I) and (II) coincide as both amount to a sign flip with probability p. Depending on the magnitude of p, the corresponding value $\lambda = \lambda_{b,p}$ may even be negative, which is unlike the case of additive noise. Recall that the error bound (12) requires $\lambda > 0$. Borrowing terminology from robust statistics, we consider $\bar{p}_b = \min\{p : \lambda_{b,p} \leq 0\}$ as the *breakdown point*, i.e., the (expected) proportion of contaminated observations that can still be tolerated so that (12) continues to hold. Mechanism (II) produces a natural counterpart to gross corruptions when the linear measurements are not subject to quantization. It is not hard to see that among all maps $\mathcal{M}^{\pm} \to \mathcal{M}^{\pm}$ applied randomly to the observations with a fixed probability, (II) maximizes the ratio $\Psi_{b,p}/\lambda_{b,p}$, hence the attribute "adversarial". In Figure 2 we display $\Psi_{b,p}/\lambda_{b,p}$ for both (I) and (II) and $b \in \{1,2,3,4\}$. Table 2 provides the corresponding breakdown points. For simplicity, $(\mathbf{t}, \boldsymbol{\mu})$ are not optimized but set to the optimal (in the sense of Lloyd-Max) choice $(\mathbf{t}_0^*, \boldsymbol{\mu}_0^*)$ in the noiseless case. The underlying derivations can be found in Appendix K.

(I)	$b=1$	$b=2$	$b=3$	$b=4$	(II)	$b=1$	$b=2$	$b=3$	$b=4$
$\bar{p}_b$	1/2	3/4	7/8	15/16	$\bar{p}_b$	1/2	0.42	0.36	0.31

Table 2: Breakdown points $\bar{p}_b$ for mechanisms (I, left half) and (II, right half) and $b \in \{1,2,3,4\}$.

Figure 2 and Table 2 provide one more argument in favour of one-bit measurements as they offer better robustness vis-à-vis adversarial corruptions. In fact, once the fraction of such corruptions reaches 0.2, $b=1$ performs best − on the measurement scale. For the milder corruption scheme

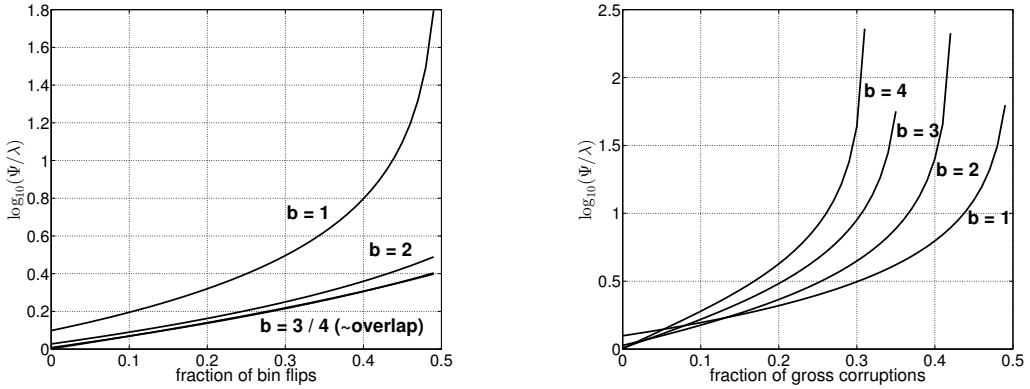

Figure 2: $\Psi_{b,p}/\lambda_{b,p}$ (on a $\log_{10}$-scale), $b \in \{1,2,3,4\}$, $p \in [0, 0.5]$ for mechanisms (I, left) and (II, right). For (I), the trajectories for $b = 3$ and $b = 4$ are indistinguishable.

(I), $b = 2$ turns out to the best choice for significant but moderate p.

4 Scale estimation

In Section 2, we have decomposed $x^* = x_u^* \psi^*$ into a product of a unit vector x_u^* and a scale parameter $\psi^* > 0$. We have pointed out that x_u^* can be estimated from a linear map of y separately from ψ^* since the latter can be absorbed into the definition of the bins $\{\mathcal{R}_k\}$. Accordingly, we may estimate x^* as $\widehat{x}\widehat{\psi}$ with $\widehat{x}$ as in (4) estimating x_u^* and $\widehat{\psi}$ estimating ψ^*. We here consider the maximum likelihood estimator (MLE) for ψ^*.

Noiseless case. To begin with, let us consider the case $\sigma = 0$, so that the $\{y_i\}$ are i.i.d. $N(0, (\psi^*)^2)$. The likelihood function is then given by

$$L(\psi) = \prod_{i=1}^{m} \sum_{k=1}^{K} I(y_i \in \mathcal{R}_k)\, \mathbf{P}(|y_i| \in \mathcal{R}_k) = \prod_{k=1}^{K} \{2(\Phi(t_k/\psi) - \Phi(t_{k-1}/\psi))\}^{m_k}, \tag{18}$$

where $m_k = |\{i : |y_i| \in \mathcal{R}_k\}|$, $k \in [K]$, and Φ denotes the standard Gaussian cdf. Note that for $K = 1$, $L(\psi)$ is constant (i.e., does not depend on ψ) which confirms that for $b = 1$, it is impossible to recover ψ^*. For $K = 2$ (i.e., $b = 2$), the MLE has a simple a closed form expression given by $\widehat{\psi} = t_1/\Phi^{-1}(0.5(1 + m_1/m))$. The following tail bound establishes fast convergence of $\widehat{\psi}$ to ψ^*.

Proposition 1. *Let $\varepsilon \in (0,1)$ and $c = 2\{\phi'(t_1/\psi^*)\}^2$, where ϕ denotes the standard Gaussian pdf and ϕ' its derivative. With probability at least $1 - 2\exp(-cm\varepsilon^2)$, we have $|\widehat{\psi}/\psi^* - 1| \le \varepsilon$.*

The exponent c is maximized for $t_1 = \psi^*$ and becomes smaller as t_1/ψ^* moves away from 1. While scale estimation from 2-bit measurements is possible, convergence can be slow if t_1/ψ^* is too small or too large.

One may profit from taking $b \ge 3$ as this introduces multiple thresholds $t_1, \ldots, t_K$, so that scale estimation is less dependent on the choice of each individual threshold. Moreover, if the thresholds are well chosen, convergence becomes faster as b increases; see [27] for a detailed analysis regarding this aspect. On the other hand, for $b \ge 3$, the MLE is no longer available in closed from.

Additive noise. We now turn to the case $\sigma > 0$. Since the $\{y_i\}$ are now i.i.d. $\sim N(0, (\psi^*)^2 + \sigma^2)$, the MLE based on (18) systematically over-estimates ψ^*. Therefore, a different approach is needed.

Suppose that x_u^* is known and denote by $[l_i, u_i] \in \{-\mathcal{R}_k\}_{k=1}^K \cup \{\mathcal{R}_k\}_{k=1}^K$ the interval the i-th observation is contained in before quantization, $i \in [m]$. Then the joint likelihood for (ψ^*, σ) is given by

$$L(\psi, \widetilde{\sigma}) = \prod_{i=1}^{m} \left\{ \Phi\left(\frac{u_i - \psi \langle a_i, x_u^* \rangle}{\widetilde{\sigma}} \right) - \Phi\left(\frac{l_i - \psi \langle a_i, x_u^* \rangle}{\widetilde{\sigma}} \right) \right\}. \tag{19}$$

Existence and uniqueness of the MLE. It is easy to see that as $\psi \to \infty$ or $\widetilde{\sigma} \to \infty$

$$\Phi\left(\frac{u_i - \psi \langle a_i, x_u^* \rangle}{\widetilde{\sigma}} \right) \to 0, \qquad \Phi\left(\frac{l_i - \psi \langle a_i, x_u^* \rangle}{\widetilde{\sigma}} \right) \to 0, \quad i \in [m],$$

and thus $L(\psi, \widetilde{\sigma}) \to 0$. Hence, the negative log-likelihood is coercive so that the MLE always exists. However, the MLE is not necessarily unique. In fact, if there exists ψ so that $\psi \langle a_i, x_u^* \rangle \in [l_i, u_i]$, $i \in [m]$ (i.e., the bin assignment of the "linear predictions" $\{\psi \langle a_i, x_u^* \rangle\}_{i=1}^m$ perfectly matches that of the observations), the likelihood (19) attains the maximum possible value of one by setting $\widetilde{\sigma}$ to zero. There are potentially multiple values of ψ satisfying the above condition; any of these is a MLE. However, this is a rather unlikely scenario as long as there is a noticeable noise level.

Computation of the MLE. A straightforward strategy to deal with the resulting two-dimensional optimization problem is coordinate descent, i.e., ψ is optimized for fixed $\widetilde{\sigma}$ and vice versa. Both these sub-problems are smooth univariate optimization problems which can be solved using methods based on golden section search as implemented in common software packages. Alternating between the two univariate problems until the objective cannot be further decreased yields a root of the likelihood equation $\nabla L(\psi, \widetilde{\sigma}) = 0$. It is known that the negative logarithm of L is convex in ψ [45], but it is not clear to us whether it is jointly convex in $(\psi, \widetilde{\sigma})$. Hence, the above likelihood equation may have multiple roots in general. Empirically, we have not encountered any issue with spurious solutions when using $\psi = 0$ and $\widetilde{\sigma}$ as the MLE from the noiseless case as starting point. Moreover, since the optimization problem is only two-dimensional, it is feasible to perform a grid search to locate a smaller region within which the MLE resides, which reduces the chance of getting trapped in an undesired stationary point different from the MLE.

So far, we have assumed that x_u^* is known. We may follow the plug-in principle and replace x_u^* by its estimator $\widehat{x}$. The resulting estimator for (ψ^*, σ) is no longer an MLE, but it can be a reasonable proxy depending on the distance of $\widehat{x}$ and x_u^*. The empirical performance of this plug-in approach is discussed in the next section.

Error bound when combining estimates for direction and scale. Suppose that $\|\widehat{x} - x_u^*\|_2 \le \delta$ and $|\widehat{\psi} - \psi^*| \le \epsilon\psi^*$. Combining these bounds, we obtain that

$$\|x^* - \widehat{x}\widehat{\psi}\|_2 = \|x_u^*\psi^* - \widehat{x}(\psi^* + (\widehat{\psi} - \psi^*))\|_2 \le \psi^*\|x_u^* - \widehat{x}\|_2 + |\widehat{\psi} - \psi^*|\|\widehat{x}\|_2 \le \psi^*(\delta + \epsilon).$$

In the noiseless case, ϵ scales as $O(1/\sqrt{m})$ in light of Proposition 1. With δ scaling as $O(w(\overline{\Delta}(\mathcal{C}))/\sqrt{m})$ according to Theorem 1, ϵ can be considered as a lower order term. The total error for estimating x^* is hence proportional to the error for estimating the direction x_u^*. Consequently, as far as the optimal trade-off between m and b is concerned, the conclusions for estimating the direction in §3.5 also apply when combining the estimates for direction and scale. In particular, since $b = 1$ renders scale estimation impossible, $b = 2$ becomes the optimal choice for a given budget of bits.

5 Experiments

We here provide numerical results supporting/illustrating some of the key points made in the previous sections. Specifically, we compare the ℓ_2-error of the estimator $\widehat{x}$ when using one-bit respectively two-bit measurements in light of what is predicted in Table 1 for different classes of signals (sparse, fused sparse, group sparse and low rank matrices) and different noise models (additive noise, random bin flips and adversarial bin flips). The empirical performance of the approach to scale estimation is also investigated. A separate set of experiments is dedicated to a comparison of the estimator studied in detail herein and possible alternatives for sparse signal recovery from quantized measurements.

5.1 One-bit vs. two-bit measurements

Setup I: additive noise. The majority of our simulations follow the model given by (1),(2) with $n = 500$, $\sigma \in \{0, 1, 2\}$ and $b \in \{1, 2\}$. For $b = 2$, quantization is performed according to the Lloyd-Max problem (16) for a $N(0, 1)$-random variable. The number of measurements m and the generation of x^* varies with $\mathcal{K}$.

1) $\mathcal{K} = B_0(s; n)$ for $s \in \{10, 20, \dots, 50\}$. The support of x^* and its signs are selected uniformly at random, while the absolute magnitude of the entries corresponding to the support are drawn from the uniform distribution on $[\beta, 2\beta]$ (for short, we write $U([\beta, 2\beta])$ in the sequel), where

$$\beta = \frac{2f}{\lambda_{1,\sigma}}\sqrt{\frac{\log n}{m}}, \quad m = \left(\frac{3f}{\lambda_{1,\sigma}}\right)^2 s \log n, \quad \text{and } f \in \{0.5, 1, 1.5, \dots, 4\}, \tag{20}$$

with f controlling signal strength. The resulting signal is then normalized to unit 2-norm.

2) Fused Sparsity (cf. §3.3) with s equally sized blocks, $s \in \{5, 10, 20, 25, 50\}$. The entries corresponding to each block are set to ± 1 in an alternating fashion, and scaled such that x^* has unit 2-norm. The number of measurements m are chosen according to (20).

3) Group Sparsity (cf. §3.3) with $L = 100$ groups and $s \in \{2, 4, 6, 8, 10\}$. For $1 \le \ell \le s$, we let

$$x^*_{G_\ell} \sim U([\beta, 2\beta]) \cdot N(0, I_{n/L}), \quad \beta = \frac{2f}{\lambda_{1,\sigma}}\sqrt{\frac{\log L}{m}}, \quad m = \left(\frac{3f}{\lambda_{1,\sigma}}\right)^2 s(n/L + \log(L)),$$

with f as in (20). We let $x^*_{G_\ell} = 0$ for $\ell > s$ and subsequently normalize x^* to unit 2-norm.

4) Low-rank matrices (cf. §3.3) with $n_1 = 50$, $n_2 = 30$, $s \in \{2, 4, 6, 8, 10\}$. We draw a random $n_1 \times n_2$-matrix G with i.i.d. $N(0, 1)$-entries and compute its SVD $G = U\Sigma V^\top$. We then let

$$X^* = U_s \mathrm{diag}(d_1, \dots, d_s) V_s^\top, \quad \{d_j\}_{j=1}^s \overset{\text{i.i.d.}}{\sim} U([\beta, 2\beta]), \quad \beta = 2f/\lambda_{1,\sigma}, \quad m = (3f/\lambda_{1,\sigma})^2 s(n_1 + n_2),$$

where U_s and V_s contain the first s columns respectively rows of U respectively V. Subsequently, X^* is normalized to unit Frobenius norm.

Each possible configuration for s, f and σ is replicated 20 times. It is straightforward to compute $\widehat{x}$ for 1)–4) even though $\mathcal{K}$ is non-convex in all cases; cf. §3.3 and Appendix I.

Setup II: beyond additive noise. In another set of simulations, we consider random and adversarial bin flips as discussed at the end of Section 3. We here only consider $\mathcal{K} = B_0(s; n)$. The setup is as in I,1) above with $\sigma = 0$ apart from the following modifications: after quantization, the respective bin flip mechanism is applied to the observations with probability p, where $p \in \{0.05, 0.1, \dots, 0.4\}$. Regarding x^* and m, we follow the same scheme as for additive noise, but replace $\lambda_{1,\sigma}$ by $\lambda_{1,p}$, cf. Appendix K.

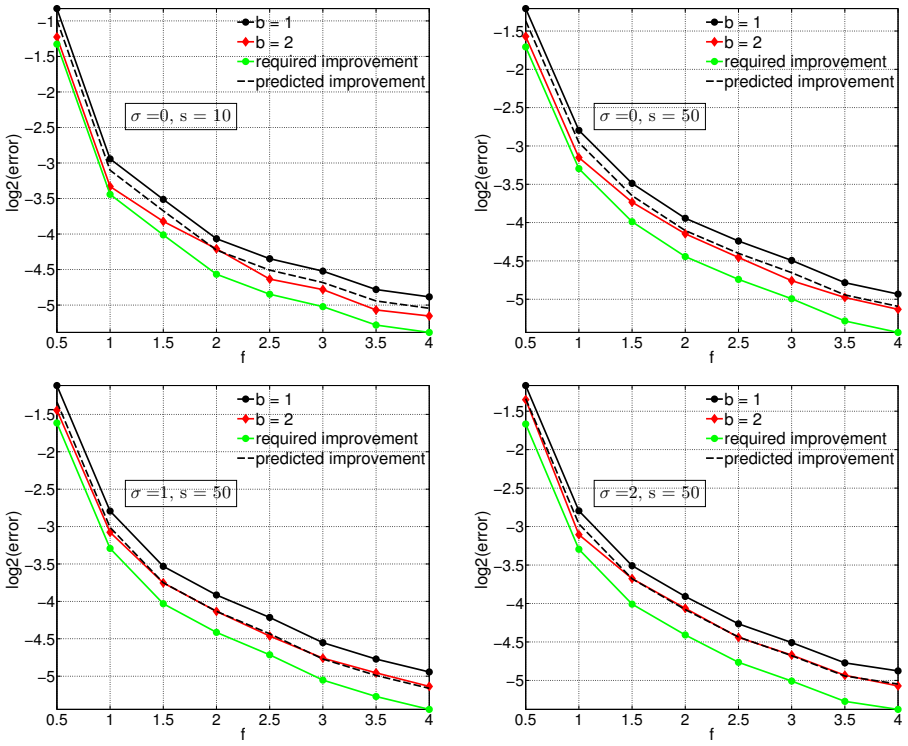

Figure 3: Average ℓ_2-estimation errors $\|x^* - \widehat{x}\|_2$ for $\mathcal{K} = B_0(s; n)$ under **additive noise** (Setup I) for $b = 1$ and $b = 2$ on the $\log_2$-scale in dependence of the signal strength f. The curve "predicted improvement" (of $b = 2$ vs. $b = 1$) is obtained by scaling the error of $b = 1$ by the factor predicted by the theory of Section 3 (cf. Table 1). Likewise the curve "required improvement" results by scaling the error of $b = 1$ by $1/\sqrt{2}$ and indicates what would be required by $b = 2$ to improve over $b = 1$ at the level of total #bits.

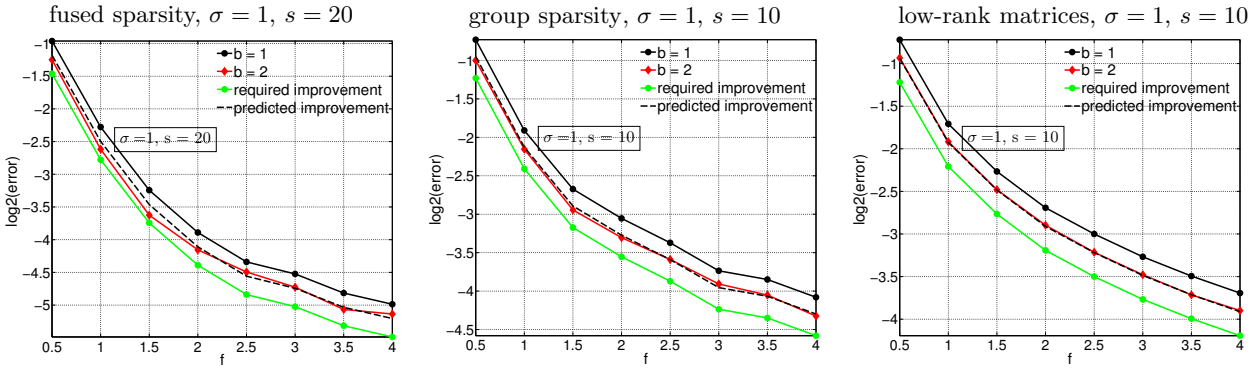
fused sparsity, σ = 1, s = 20
group sparsity, σ = 1, s = 10
low-rank matrices, σ = 1, s = 10
b = 1
b = 2
required improvement
predicted improvement
σ =1, s = 20
σ =1, s = 10
log2(error)
f

Figure 4: Average ℓ_2-estimation errors $\|x^* - \widehat{x}\|_2$ for the other three classes of signals 2)–4). The annotation of the plots follows that in Figure 3.

Conclusion. The experiments reveal that what is predicted by the analysis concerning the relative performance of 1-bit and 2-bit measurements for estimating x^* closely agrees with what is observed empirically, as can be seen in Figures 3, 4, 5 and 6. The agreement is consistent over different signal and noise models.

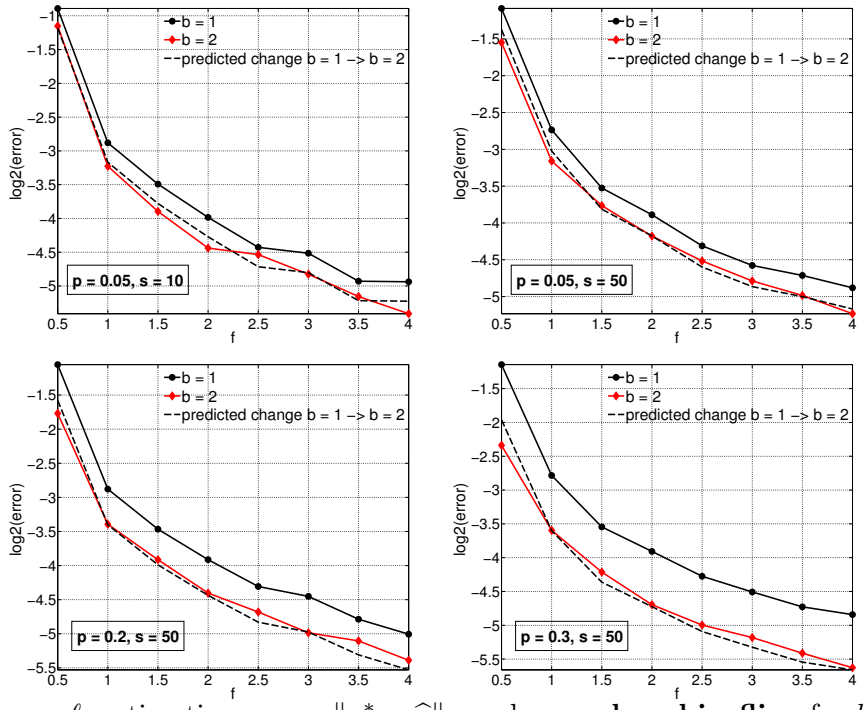

Figure 5: Average ℓ_2-estimation errors $\|x^* - \widehat{x}\|_2$ under **random bin flips** for $b = 1$ and $b = 2$ on the $\log_2$-scale in dependence of the signal strength f. The curve "predicted change $1 \to 2$" is obtained by scaling the error of $b = 1$ by the factor predicted by the theory of Section 3.

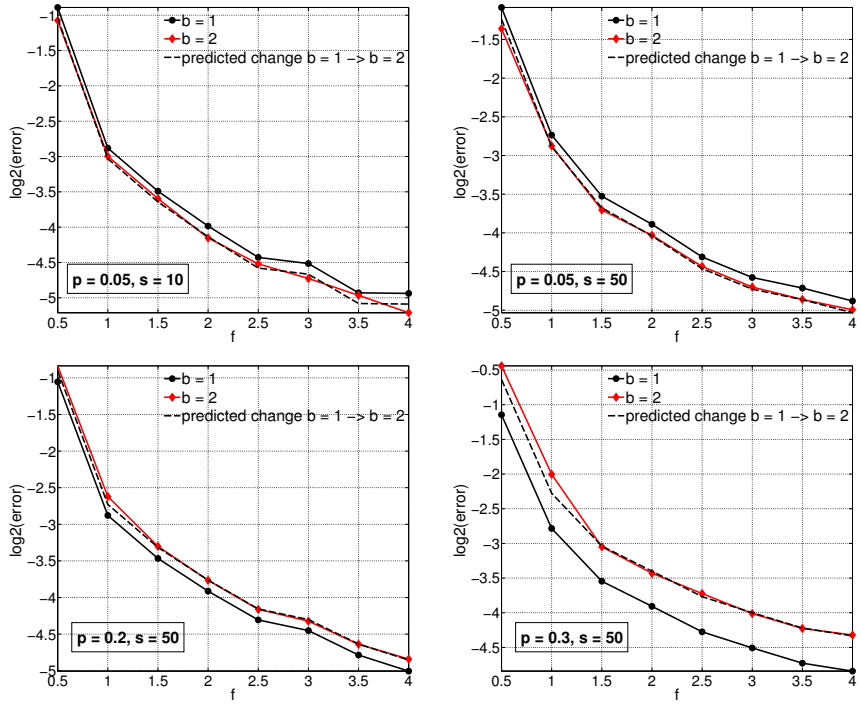

Figure 6: Average ℓ_2-estimation errors $\|x^* - \widehat{x}\|_2$ under **adversarial bin flips** for $b = 1$ and $b = 2$ on the $\log_2$-scale in dependence of the signal strength f. The curve "predicted change $1 \to 2$" is obtained by scaling the error of $b = 1$ by the factor predicted by the theory of Section 3.

5.2 Estimation of the scale and the noise level

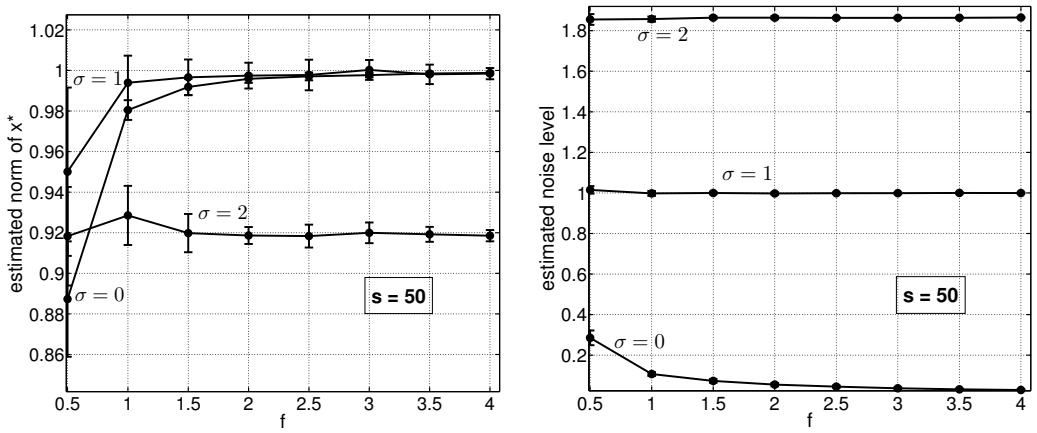

Figure 7: Estimation of $\psi = \|x^*\|_2$ (here 1) and σ. The curves depict the average of the plug-in MLE discussed in Section 4 while the bars indicate ± 1 standard deviation.

Figure 7 suggests that the plug-in MLE for $(\psi^* = \|x^*\|_2, \sigma)$ as outlined in Section 4 is a suitable approach, at least as long as ψ^*/σ is not too small. For $\sigma = 2$, the plug-in MLE for ψ^* appears to have a noticeable bias as it tends to 0.92 instead of 1 for increasing f (and thus increasing m). Observe that for $\sigma = 0$, convergence to the true value $\psi^* = 1$ is slower as for $\sigma = 1$, while σ is over-estimated (about 0.2) for small f. The above two issues are presumably a plug-in effect, i.e., a consequence of using $\widehat{x}$ in place of x^*_u.

5.3 Alternative recovery algorithms

For $\mathcal{K} = B_0(s; n)$, we compare the empirical performance of the linear estimator (4) to several alternatives. Two of those are based on a more principled albeit computationally more involved approach that tries to enforce agreement of $Q(y)$ and $Q(A\widehat{x})$ w.r.t. the Hamming distance (or a surrogate thereof), thereby using knowledge about the quantization map unlike (4). One may thus expect that the performance of the latter is inferior. In summary, our experiments confirm that this is true in low-noise settings, but not so if the noise level is substantial. Below we briefly present the alternatives that we consider.

Plan-Vershynin: The approach in [33] that differs from (4) only in that it uses a convex relaxation of the constraint set of the form $\mathcal{C} = B_0(s; n) \cap B_2^n \subset \sqrt{s} B_1^n \cap B_2^n$. It thus falls under the framework of linear signal recovery outlined in §2, and its inclusion is mainly for the sake of reference given the popularity of the work [33]. As shown in Figure 8 the performance is similar though slightly inferior to (4).

IHT-quadratic: The standard Iterative Hard Thresholding algorithm based on quadratic loss [5]. When using quadratic loss, one does not take into account the fact that the observations are quantized. The estimator (4) can be seen as one-step version of Iterative Hard Thresholding.

Lasso: ℓ_1-regularized least squares [39] according to the optimization problem $\min_x \|y - Ax\|_2^2 + \lambda \|x\|_1$. The parameter λ is selected based on five-fold cross-validation over the grid $\sqrt{m \log n} \cdot \{2^{-5}, 2^{-4}, \ldots, 2^2\}$.

The competitors listed below do no longer fit into the framework of linear recovery but into the more sophisticated class described in the introductory portion of this paragraph.

IHT-hinge ($b = 1$): The variant of Iterative Hard Thresholding for binary observations using a hinge loss-type loss function as proposed in [22].

IHT-Jacques ($b = 2$): A variant of Iterative Hard Thresholding for quantized observations based on a specific piecewise linear loss function as suggested in a paper of Jacques and collaborators [21].

SVM ($b = 1$): Linear SVM with squared hinge loss and an ℓ_1-penalty on the weights as implemented in `LIBLINEAR` [15]. The cost parameter is chosen by means of five-fold cross-validation over the grid $\sqrt{1/(m \log m)} \cdot \{2^{-3}, 2^{-2}, \ldots, 2^3\}$.

SVM-type ($b = 2$): This approach is based on the following convex optimization problem:

$$\min_{x,\{\xi_i\}} \gamma \|x\|_1 + \sum_{i=1}^{m} \xi_i \quad \text{subject to} \quad l_i - \xi_i \leq \langle a_i, x \rangle \leq u_i + \xi_i, \quad \xi_i \geq 0, \quad i \in [m],$$

where $[l_i, u_i]$ is the bin the i-th observations is assigned to. The essential idea is to enforce consistency of the observed and predicted bin assignments up to slacks $\{\xi_i\}$ while promoting sparsity of the solution by means of an ℓ_1-penalty. The parameter γ is chosen via five-fold cross-validation over the grid $\sqrt{m \log m} \cdot \{2^{-10}, 2^{-9}, \ldots, 2^3\}$.

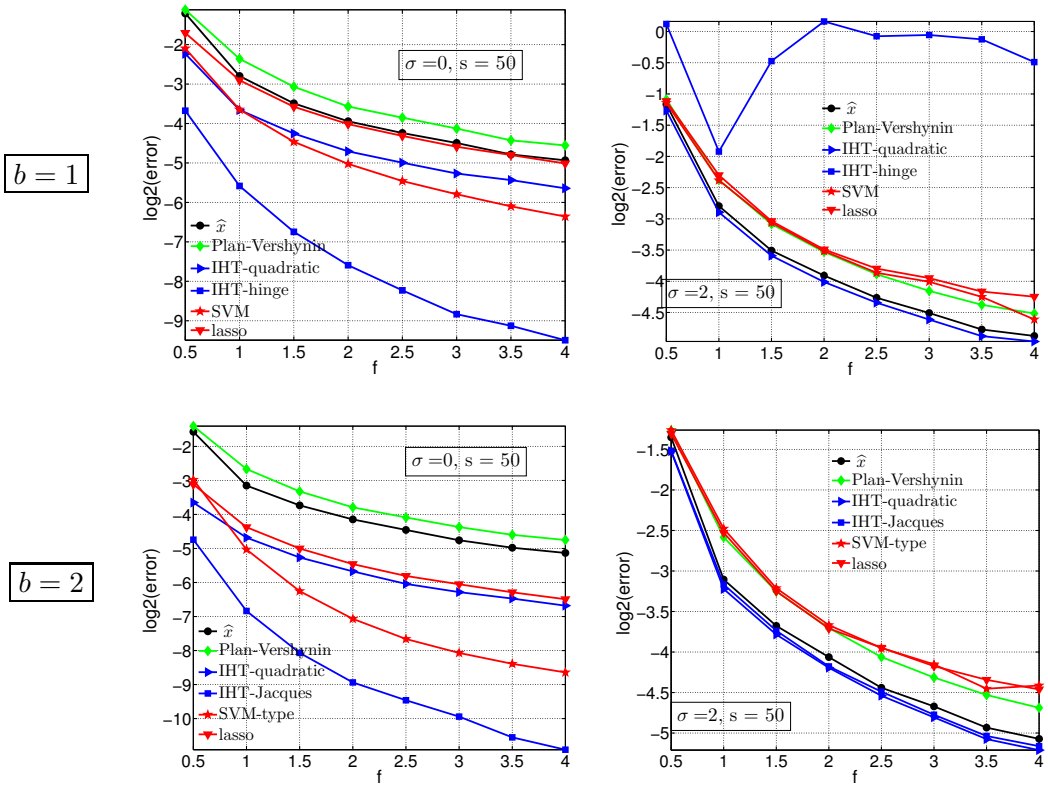

Figure 8: Average ℓ_2-estimation errors of several recovery algorithms on the $\log_2$-scale in dependence of the signal strength f. We contrast $\sigma = 0$ (L) vs. $\sigma = 2$ (R), $b = 1$ (T) vs. $b = 2$ (B).

Turning to the results as depicted by Figure 8, the difference between a noiseless ($\sigma = 0$) and heavily noisy setting ($\sigma = 2$) is perhaps most striking.

$\sigma = 0$: IHT-hinge respectively IHT-Jacques and SVM-(type) significantly outperform $\widehat{x}$. By comparing the errors for IHT and SVM, $b = 2$ can be seen to improve over $b = 1$ at the level of the total # bits.

$\sigma = 2$: the canonical linear estimator is on par with the best performing methods. IHT-quadratic

for $b = 2$ only achieves a moderate reduction in error over $b = 1$, while IHT-hinge is supposedly affected by convergence issues that are known to arise in a high-noise setting [19].

Overall, the results suggest that a setting with substantial noise favours a crude approach (low-bit measurements and conceptually simple recovery algorithms).

6 Conclusion

We have considered linear signal recovery from b-bit quantized measurements. The main finding is that increasing the number of bits per measurement beyond is not efficient at the bit scale since the reduction in estimation error at the measurement scale is not significant enough. A compelling argument in favor of two-bit measurements is the fact that the norm of the signal can be estimated unlike the case one-bit measurements. Compared to high-precision measurements, two-bit measurements also exhibit strong robustness properties. It is of interest if and under what circumstances the conclusion may differ for other recovery algorithms. Experimental results suggest a transition between settings with low noise on the one hand and substantial noise on the other hand. For example, in a setting without noise we have seen improvements at the level of bits for iterative hard thresholding algorithms when increasing b from one to two. Providing theoretical support for this observation constitutes an interesting topic for future research.

Acknowledgement

The work is partially supported by NSF-Bigdata-1419210, NSF-III-1360971, ONR-N00014-13-1-0764, and AFOSR-FA9550-13-1-0137.

Appendix: Proofs and Derivations

A Linearity lemma

The following simple result is fundamental to linear signal recovery. The lemma implies that for the non-linear model (1), it may still be possible to recover the signal by a linear approach depending on the nature of the nonlinear map Q. Let us recall the definition of the map θ in (8) according to $\mathbf{E}[y_1|a_1] = \theta(\langle a_1, x_u^*\rangle)$, $\lambda = \mathbf{E}[g\,\theta(g)]$ with $g \sim N(0,1)$, and $\eta = A^\top y/m$.

Lemma A.1. *For all $x \in \mathbb{R}^n$, we have $\mathbf{E}[\langle x, \eta\rangle] = \lambda \langle x, x_u^*\rangle$. In particular, by considering $x = e_j$, $j \in [n]$, where $\{e_j\}_{j=1}^n$ is the standard basis of $\mathbb{R}^n$, we have $\mathbf{E}[\eta] = \lambda x_u^*$.*

Proof. The proof below is taken from [33], Lemma 4.1. We include it here to make the paper self-contained. We have

$$\begin{aligned}
\mathbf{E}[\langle x, \eta\rangle] &= \mathbf{E}\left[\left\langle x, A^\top y/m\right\rangle\right] \\
&= \mathbf{E}[\langle Ax, y\rangle /m] \\
&= \mathbf{E}[\langle a_1, x\rangle\, y_1] \\
&= \mathbf{E}\,\mathbf{E}[y_1 \langle a_1, x\rangle |a_1] \\
&= \mathbf{E}[\theta(\langle a_1, x_u^*\rangle) \langle a_1, x\rangle] \\
&= \mathbf{E}\left[\theta(\langle a_1, x_u^*\rangle) \left\langle a_1, x^\parallel + x^\perp\right\rangle\right] \\
&= \langle x, x_u^*\rangle\, \mathbf{E}[\theta(g)g],\ g \sim N(0,1) \\
&= \lambda \langle x, x_u^*\rangle,
\end{aligned}$$

where in the third line from the bottom $x^\parallel = \langle x, x_u^*\rangle\, x_u^*$ and $x^\perp$ denote the orthogonal projection of x on x_u^* and its orthogonal complement, respectively. We then use that $\left\langle a_1, x^\perp\right\rangle$ and $\langle a_1, x_u^*\rangle$ are Gaussian and uncorrelated and hence also independent random variables. □

Note that λ can be zero (e.g., if θ is even). This is not a concern if model (1) holds and Q is the scalar b-bit quantizer in (2) (cf. Lemma 1). For the alternative noise models in §3.5, the situation becomes different as is explained in detail in Appendix K.

B Proof of Theorem 1

The proof of our result relies on the previous Lemma and a series of results from the theory of empirical processes which are relegated to separate sections in this appendix. The structure largely builds on techniques used in [33].

Step 1. Basic inequality.

Since $\widehat{x}$ in (4) is a minimizer and x_u^* is a feasible solution, we have

$$-\langle \eta, \widehat{x}\rangle \le -\langle \eta, x_u^*\rangle .$$

Subtracting and adding $\mathbf{E}[\eta]$, we obtain after some re-arrangements that

$$\langle x_u^* - \widehat{x}, \mathbf{E}[\eta]\rangle \le \langle \widehat{x} - x_u^*, \eta - \mathbf{E}[\eta]\rangle . \tag{21}$$

Step 2. Lower bounding the left hand side.

From Lemma A.1, we obtain that $\mathbf{E}[\eta] = \lambda x_u^*$. Therefore,

$$\langle x_u^* - \widehat{x}, \mathbf{E}[\eta] \rangle = \lambda \langle x_u^* - \widehat{x}, x_u^* \rangle \tag{22}$$

On the other hand, we have that

$$\|x_u^* - \widehat{x}\|_2^2 \leq 2(1 - \langle \widehat{x}, x_u^* \rangle) = 2(\langle x_u^*, x_u^* \rangle - \langle \widehat{x}, x_u^* \rangle) = 2(\langle x_u^*, x_u^* - \widehat{x} \rangle), \tag{23}$$

using that $\|x_u^*\|_2 = 1$ and $\|\widehat{x}\|_2 \leq 1$. Combining (22) and (23), the l.h.s. of (21) is bounded as

$$\langle x_u^* - \widehat{x}, \mathbf{E}[\eta] \rangle \geq \frac{\lambda}{2} \|x_u^* - \widehat{x}\|_2^2. \tag{24}$$

Step 3.1 Upper bounding the right hand side.
Recalling the definition of $\overline{\Delta}(\mathcal{C}) = \overline{\Delta}(\mathcal{C}; x_u^*)$ in (10), we have that

$$\langle \widehat{x} - x_u^*, \eta - \mathbf{E}[\eta] \rangle \leq \sup_{v \in \overline{\Delta}(\mathcal{C})} \langle v, \eta - \mathbf{E}[\eta] \rangle \ \|\widehat{x} - x_u^*\|_2. \tag{25}$$

The supremum on the right hand sided is controlled by first bounding its expectation and then establishing concentration.

Step 3.2 Upper bounding the expectation.
In the sequel, $\{\xi_i\}_{i=1}^n$ denote i.i.d. Rademacher variables taking values in $\{\pm 1\}$ with equal probability. Invoking Lemma L.1 with $\Gamma = \{|\langle v, \cdot \rangle|, \ v \in \overline{\Delta}(\mathcal{C})\}$, we obtain that

$$\begin{aligned}
\mathbf{E}\left[\sup_{v \in \overline{\Delta}(\mathcal{C})} |\langle \eta - \mathbf{E}[\eta], v \rangle|\right] &= \mathbf{E}\left[\sup_{v \in \overline{\Delta}(\mathcal{C})} \left|\frac{1}{m}\sum_{i=1}^m (y_i \langle a_i, v \rangle - \mathbf{E}[y_i \langle a_i, v \rangle])\right|\right] \\
&\overset{\text{Lemma L.1}}{\leq} 2\,\mathbf{E}\left[\sup_{v \in \overline{\Delta}(\mathcal{C})} \left|\frac{1}{m}\sum_{i=1}^m y_i \langle \widetilde{a}_i, v \rangle\right|\right], \quad \widetilde{a}_i := a_i \xi_i, \ i = 1, \ldots, m \\
&= 2\,\mathbf{E}_y\left[\mathbf{E}\left[\sup_{v \in \overline{\Delta}(\mathcal{C})} \left|\frac{1}{m}\left\langle \sum_{i=1}^m y_i \widetilde{a}_i, v \right\rangle\right| \,\middle|\, y\right]\right] \\
&\overset{\mathbf{(C)}}{\underset{n \to \infty}{=}} 2\,\mathbf{E}_y\left[\mathbf{E}\left[\sup_{v \in \overline{\Delta}(\mathcal{C})} \left|\frac{\|y\|_2}{m} \langle g, v \rangle\right|\right]\right], \quad g \sim N(0, I_n), \\
&= \frac{2\,\mathbf{E}[\|y\|_2] w(\overline{\Delta}(\mathcal{C}))}{m} \\
&\leq \frac{2\sqrt{\mathbf{E}[\|y\|_2^2]} w(\overline{\Delta}(\mathcal{C}))}{m} \\
&= \frac{2\Psi_{b,\sigma}(\mathbf{t}, \boldsymbol{\mu})\, w(\overline{\Delta}(\mathcal{C}))}{\sqrt{m}}.
\end{aligned} \tag{26}$$

The fourth line from the top holds under condition **(C)** as $n \to \infty$ because a_i and y_i are asymptotically independent, $i \in [m]$. In fact, for any $j \in [n]$, we have

$$\mathbf{E}[\langle a_i, x_u^* \rangle a_{ij}] = \mathbf{E}[\langle a_i, x_u^* \rangle \langle a_i, e_j \rangle] = (x_u^*)_j \to 0 \quad \text{as } n \to \infty.$$

Since $\langle a_i, x_u^* \rangle$ and a_{ij} are jointly Gaussian with vanishing covariance, they are also asymptotically independent. As y_i results as a transformation of $\langle a_i, x_u^* \rangle$ and ε_i, it holds that a_{ij} and y_i are asymptotically independent as well, $j \in [n]$, $i \in [m]$. The same is true for $\widetilde{a}_{ij}$ and y_i, $j \in [n]$, $i \in [m]$. In

this situation, conditional on y, $\sum_{i=1}^m y_i a_{ij}$ follows a $N(0, \|y\|_2^2)$-distribution, $j \in [n]$. The remaining lines follow from the rotational invariance of the Gaussian distribution, the definitions of $\Psi_{b,\sigma}(\mathbf{t}, \boldsymbol{\mu})$ and $w(\overline{\Delta}(\mathcal{C}))$, and Jensen's inequality.

Step 3.3 Concentration.
Using Lemma L.2 with $\Gamma = \{|\langle v, \cdot \rangle|,\ v \in \overline{\Delta}(\mathcal{C})\}$ and Lemma L.3, we establish that the bound (26) above gives rise to an upper bound for the empirical process $\{\langle v, \eta - \mathbf{E}[\eta] \rangle,\ v \in \overline{\Delta}(\mathcal{C})\}$ in (21) that holds with high probability. We start by upper bounding the right hand side of Lemma L.2, following the strategy of the previous step, i.e., we condition on y and then use asymptotic independence. For any $u > 0$

$$\mathbf{P}\left(\sup_{v \in \overline{\Delta}(\mathcal{C})} \left| \frac{1}{m} \sum_{i=1}^m \xi_i y_i \langle a_i, v \rangle \right| > u/2 \right) = \mathbf{E}_y \left[\mathbf{P}\left(\sup_{v \in \overline{\Delta}(\mathcal{C})} \left| \frac{1}{m} \sum_{i=1}^m \xi_i y_i \langle a_i, v \rangle \right| > u/2 \,\middle|\, y \right) \right],$$
$$\overset{\mathbf{(C)}}{\underset{n \to \infty}{=}} \mathbf{E}_y \left[\mathbf{P}\left(\frac{\|y\|}{m} \sup_{v \in \overline{\Delta}(\mathcal{C})} |\langle g, v \rangle| > u/2 \right) \right]. \tag{27}$$

Consider the event $\{m^{-1/2}\|y\|_2 > 2\Psi_{b,\sigma}(\mathbf{t}, \boldsymbol{\mu})\}$. Note that the $\{y_i\}_{i=1}^m$ are i.i.d. zero-mean sub-Gaussian random variables with standard deviation $\Psi_{b,\sigma}(\mathbf{t}, \boldsymbol{\mu})$. Then by Lemma L.4 and Jensen's inequality,

$$\mathbf{P}(m^{-1/2}\|y\|_2 > 2\Psi_{b,\sigma}(\mathbf{t}, \boldsymbol{\mu})) \leq \exp(-cm), \tag{28}$$

for some universal constant $c > 0$. We work thus conditional on the event $E_0 = \{m^{-1/2}\|y\|_2 \leq 2\Psi_{b,\sigma}(\mathbf{t}, \boldsymbol{\mu})\}$. Accordingly,

$$\mathbf{E}_y \left[\mathbf{P}\left(\frac{\|y\|}{m} \sup_{v \in \overline{\Delta}(\mathcal{C})} |\langle g, v \rangle| > u/2 \right) \right] \leq \mathbf{P}\left(\frac{2\Psi_{b,\sigma}(\mathbf{t}, \boldsymbol{\mu})}{\sqrt{m}} \sup_{v \in \overline{\Delta}(\mathcal{C})} |\langle g, v \rangle| > u/2 \right) \mathbf{P}(E_0) + (1 - \mathbf{P}(E_0))$$
$$\leq \mathbf{P}\left(\sup_{v \in \overline{\Delta}(\mathcal{C})} |\langle g, v \rangle| > \frac{\sqrt{m} \cdot u}{4\Psi_{b,\sigma}(\mathbf{t}, \boldsymbol{\mu})} \right) + (1 - \mathbf{P}(E_0)) \tag{29}$$

Choosing $u = 8\Psi_{b,\sigma}(\mathbf{t}, \boldsymbol{\mu}) w(\overline{\Delta}(\mathcal{C}))/\sqrt{m}$, we invoke Lemma L.3 with $\mathcal{T} = \overline{\Delta}(\mathcal{C})$, $X_t = |\langle g, t \rangle|$, $t \in \mathcal{T}$, $\mathbf{E}[\sup_{t \in \mathcal{T}} X_t] = w(\overline{\Delta}(\mathcal{C}))$ and $r = w(\overline{\Delta}(\mathcal{C}))$, and combining (27), (28), (29), we obtain that

$$\mathbf{P}\left(\sup_{v \in \overline{\Delta}(\mathcal{C})} \left| \frac{1}{m} \sum_{i=1}^m \xi_i y_i \langle a_i, v \rangle \right| > \frac{8}{2} \frac{\Psi_{b,\sigma}(\mathbf{t}, \boldsymbol{\mu}) w(\overline{\Delta}(\mathcal{C}))}{\sqrt{m}} \right) \leq \exp(-w(\overline{\Delta}(\mathcal{C}))/2) + \exp(-cm). \tag{30}$$

Step 4. Putting together the pieces
Comparing (30) with the right hand side of (42) in Lemma L.2 and combining that with (26) to bound the expectation appearing in the left hand side of (42)

$$\mathbf{P}\left(\sup_{v \in \overline{\Delta}(\mathcal{C})} |\langle \eta - \mathbf{E}[\eta], v \rangle| \geq \frac{12\Psi_{b,\sigma}(\mathbf{t}, \boldsymbol{\mu}) w(\overline{\Delta}(\mathcal{C}))}{\sqrt{m}} \right) \leq 4(\exp(-w(\overline{\Delta}(\mathcal{C}))/2) + \exp(-cm)). \tag{31}$$

Combining (31) with (24) and (25), the proof is complete.

C Proof of Theorem 2

Let us denote $\zeta := \Sigma^{-1}\eta$, and $z^* := \Sigma^{1/2}x^*$. Without loss of generality, suppose that the thresholds $\mathbf{t}$ are scaled such that it is equivalent to assume that $\|z^*\|_2 = 1$, cf. the discussion following (3).

We start by noting that $\mathbf{E}[\zeta] = \lambda x^*$:

$$\begin{aligned}\mathbf{E}[\zeta] = \Sigma^{-1}\,\mathbf{E}[\eta] = \Sigma^{-1}\,\mathbf{E}\left[\sum_{i=1}^m y_i a_i/m\right] \\ = \Sigma^{-1}\,\mathbf{E}[a_1 y_1] = \Sigma^{-1/2}\,\mathbf{E}\left[a_1^0 Q(\langle a_1^0, z^*\rangle + \sigma\varepsilon_i)\right] = \Sigma^{-1/2}\lambda z^* = \lambda x^*,\end{aligned}$$

where $a_1^0 \sim N(0, I_n)$. The penultimate inequality follows from Lemma A.1. We now proceed as in the proof of Theorem 1. Since x_u^* is feasible for the optimization problem (14), we have $-\langle\zeta, \widehat{x}_\Sigma\rangle \le -\langle\zeta, x_u^*\rangle$. After some manipulations, we obtain that

$$\langle x_u^* - \widehat{x}_\Sigma, \mathbf{E}[\zeta]\rangle \le \langle \widehat{x}_\Sigma - x_u^*, \zeta - \mathbf{E}[\zeta]\rangle \tag{32}$$

The left hand side of this inequality can be lower bounded as follows:

$$\begin{aligned}\langle x_u^* - \widehat{x}_\Sigma, \mathbf{E}[\zeta]\rangle = \langle x_u^* - \widehat{x}_\Sigma, \lambda x^*\rangle &= \lambda\|x^*\|_2 - \lambda\|x^*\|_2\,\langle \widehat{x}_\Sigma, x_u^*\rangle \\ &= \lambda\|x^*\|_2(1 - \langle \widehat{x}_\Sigma, x_u^*\rangle) \\ &\ge \frac{\lambda\|x^*\|_2}{2}\|\widehat{x}_\Sigma - x_u^*\|_2^2 \\ &\ge \frac{\lambda}{2\|\Sigma^{1/2}\|}\|\widehat{x}_\Sigma - x_u^*\|_2^2. \end{aligned}\tag{33}$$

For the last line, we have used that $\|x^*\|_2 = \|\Sigma^{-1/2}z^*\|_2 \ge \|z^*\|_2/\|\Sigma^{1/2}\| = 1/\|\Sigma^{1/2}\|$, where $\|M\|$ denotes the spectral norm of a matrix M.

In order to upper bound the right hand side of inequality (32), we follow the approach in the proof of Theorem 1 up to a single modification.
A key step therein is to control $\mathbf{E}\left[\sup_{v\in\overline{\Delta}(\mathcal{C})}\left|\frac{1}{m}\sum_{i=1}^m y_i\,\langle a_i\xi_i, v\rangle\right|\right]$ in Step 3.2 above which becomes $\mathbf{E}\left[\sup_{v\in\overline{\Delta}(\mathcal{C})}\left|\frac{1}{m}\sum_{i=1}^m y_i\,\langle \Sigma^{-1}a_i\xi_i, v\rangle\right|\right]$ in the anisotropic case.

Writing $a_i = \Sigma^{1/2}a_i^0$, where $a_i^0 \sim N(0, I_n)$ and following the approach in Step 3.2 above, we end up with the problem of controlling $\mathbf{E}\left[\sup_{v\in\overline{\Delta}(\mathcal{C})}\left|\left\langle \Sigma^{-1/2}g, v\right\rangle\right|\right]$. By the Sudakov-Fernique comparison inequality (e.g., [2], Theorem 2.2.3),

$$\mathbf{E}\left[\sup_{v\in\overline{\Delta}(\mathcal{C})}\left|\left\langle \Sigma^{-1/2}g, v\right\rangle\right|\right] \le \|\Sigma^{-1/2}\|\,\mathbf{E}\left[\sup_{v\in\overline{\Delta}(\mathcal{C})}|\langle g, v\rangle|\right] = \|\Sigma^{-1/2}\|w(\overline{\Delta}(\mathcal{C})).$$

Thus, carrying an extra factor of $\|\Sigma^{-1/2}\|$, we may continue with the remaining steps in the proof of Theorem 1. Using (33) and noting that $\kappa(\Sigma) = \|\Sigma^{-1/2}\|\|\Sigma^{1/2}\|$ yields the claim.

D Asymptotic sharpness of the leading constant

We here derive (13) as $m, n \to \infty$, $n/m \to 0$ and **(C)** holds.

We here suppress dependence of λ and Ψ on $\mathbf{t}, \boldsymbol{\mu}$, b and σ. Fix $j \in [n]$ arbitrary and let $\sigma_i^2 = \mathrm{Var}[A_{ij}y_i] = \mathbf{E}[\{A_{ij}y_i - \lambda x_{u,j}^*\}^2] = \mathbf{E}[A_{ij}y_i^2] - \lambda(x_{u,j}^*)^2$, $i \in [m]$, and let $s_m^2 = \sum_{i=1}^m \sigma_i^2$. Since

the $\{A_{ij}y_i - \lambda x^*_{u,j}\}_{i=1}^m$ are independent zero-mean sub-Gaussian random variables, the Lyapunov condition

$$\frac{1}{s_m^3}\sum_{i=1}^m \mathbf{E}[|A_{ij}y_i - \lambda x^*_{u,j}|^3] \to 0 \quad \text{as } m \to \infty$$

is trivially satisfied. Consequently, by the Lyapunov central limit theorem ([4], p. 362)

$$\frac{1}{s_m}\sum_{i=1}^m (A_{ij}y_i - \lambda x^*_{u,j}) \xrightarrow{\mathcal{D}} N(0,1) \quad \text{as } m \to \infty. \tag{34}$$

Let us re-consider $\sigma_i^2 = \mathbf{E}[(A_{ij}y_i - \lambda x^*_{u,j})^2]$. As $n \to \infty$, A_{ij} and y_i become independent under **(C)** and thus $\sigma_i^2 \to \mathbf{E}[A_{ij}^2]\,\mathbf{E}[y_i^2] - \lambda(x^*_{u,j})^2 = \Psi^2 - \lambda(x^*_{u,j})^2$, $i \in [m]$. Therefore,

$$\frac{s_m^2}{m(\Psi^2 - \lambda(x^*_{u,j})^2)} \to 1 \quad \text{as } m, n \to \infty.$$

Moreover, under **(C)**, $m(\Psi^2 - \lambda(x^*_{u,j})^2) = m\Psi^2 - o(m)$ as $m, n \to \infty$ and hence also

$$\frac{s_m^2}{m\Psi^2} \to 1 \quad \text{as } m, n \to \infty. \tag{35}$$

Combining (34) and (35), it follows that

$$\sqrt{m}\left(\frac{\sum_{i=1}^m A_{ij}y_i}{m} - \lambda x^*_{u,j}\right) \xrightarrow{\mathcal{D}} N(0, \Psi^2) \quad \text{as } m, n \to \infty.$$

Now consider $\|A^\top y/m\|_2^2$. As $m, n \to \infty$ and $n/m \to 0$, this quantity converges to $\lambda^2\|x_u^*\|_2^2 = \lambda^2$ in probability. By the continuous mapping theorem and Slutsky's theorem ([40], §2), we finally obtain that

$$\sqrt{m}\left(\frac{\sum_{i=1}^m A_{ij}y_i/m}{\|A^\top y/m\|_2} - x^*_{u,j}\right) \xrightarrow{\mathcal{D}} N\left(0, \frac{\Psi^2}{\lambda^2}\right) \quad \text{as } m, n \to \infty.$$

The result to be shown follows by noting that for $\mathcal{K} = \mathbb{R}^n$, we have $\widehat{x} = \eta/\|\eta\|_2$ and accordingly $\widehat{x}_j = \frac{\sum_{i=1}^m A_{ij}y_i/m}{\|A^\top y/m\|_2}$.

E Proof of Lemma 1

The proof of Lemma 1 requires three additional lemmas.

Lemma E.1. *Let $g \sim N(0,1)$ and $\zeta : \mathbb{R} \to \mathbb{R}$ be any differentiable function satisfying $|\zeta(x)x\phi(x)| \to 0$ as $x \to \infty$, where ϕ denotes the standard Gaussian pdf. Then $\mathbf{E}[\zeta(g)g] = \mathbf{E}[\zeta'(g)]$.*

Proof. Observe that $\phi'(x) = -x\phi(x),\ x \in \mathbb{R}$. Using integration by parts we thus have

$$\begin{aligned}\mathbf{E}[\zeta(g)g] &= \int_{\mathbb{R}} x\zeta(x)\phi(x)\,dx = \big\{\zeta(x)(-\phi'(x))\big\}\Big|_{-\infty}^{\infty} + \int_{\mathbb{R}} \zeta'(x)\,\phi(x)\,dx \\ &= \int_{\mathbb{R}} \zeta'(x)\,\phi(x)\,dx = \mathbf{E}[\zeta'(g)].\end{aligned}$$

□

Lemma E.2. *For all $\alpha, \beta > 0$ and all $\mu, \nu \in \mathbb{R}$, one has*

$$\int_{-\infty}^{\infty} \frac{1}{\alpha}\phi\left(\frac{x-\mu}{\alpha}\right)\frac{1}{\beta}\phi\left(\frac{x-\nu}{\beta}\right)dx = \frac{1}{\sqrt{\beta^2+\alpha^2}}\phi\left(\frac{\mu-\nu}{\sqrt{\beta^2+\alpha^2}}\right).$$

Proof. Using elementary manipulations, one computes

$$\begin{aligned}
&\int_{-\infty}^{\infty} \frac{1}{\alpha}\phi\left(\frac{x-\mu}{\alpha}\right)\frac{1}{\beta}\phi\left(\frac{x-\nu}{\beta}\right)dx \\
&= \frac{1}{2\pi\alpha\beta}\int_{-\infty}^{\infty} \exp\left(-\frac{(\mu-x)^2}{2\alpha^2}\right)\exp\left(-\frac{(\nu-x)^2}{2\beta^2}\right)dx \\
&= \frac{1}{2\pi\alpha\beta}\exp\left(-\frac{(\mu-\nu)^2}{2(\alpha^2+\beta^2)}\right)\times \\
&\times \int_{-\infty}^{\infty} \exp\left(-\frac{1}{2}\left(\left(\frac{\mu}{\alpha^2}+\frac{\nu}{\beta^2}\right)\left(\frac{1}{\alpha^2}+\frac{1}{\beta^2}\right)-x\right)^2\left(\frac{1}{\alpha^2}+\frac{1}{\beta^2}\right)\right)dx \\
&= \frac{1}{2\pi\alpha\beta}\sqrt{2\pi}\frac{\alpha\beta}{\sqrt{\beta^2+\alpha^2}}\exp\left(-\frac{(\mu-\nu)^2}{2(\alpha^2+\beta^2)}\right) \\
&= \frac{1}{\sqrt{2\pi}}\frac{1}{\sqrt{\beta^2+\alpha^2}}\exp\left(-\frac{(\mu-\nu)^2}{2(\alpha^2+\beta^2)}\right).
\end{aligned}$$

□

Lemma E.3. *Let h be a random variable with a $N(0,\sigma^2)$-distribution. Then for any $a, b \in \mathbb{R} \cup \{-\infty, \infty\}$, $a < b$, we have*

$$\mathbf{E}[h|h \in (a,b)] = \sigma\frac{\phi(a/\sigma)-\phi(b/\sigma)}{\Phi(b/\sigma)-\Phi(a/\sigma)},$$

where Φ denotes the standard Gaussian cdf.

Proof. We have

$$\mathbf{E}[h|h \in (a,b)] = \frac{1}{\Phi(b/\sigma)-\Phi(a/\sigma)}\int_a^b \frac{x}{\sigma}\phi(x/\sigma)\,dx.$$

Using the change of variables $z = x/\sigma$ and the fact that $\phi'(z) = -z\phi(z)$, the result follows. □

Before finally turning to the proof of Lemma 1, let us recall the definition of the quantization map (2). In that definition we have used the symmetry of the Gaussian distribution around 0 so that a partitioning of $\mathbb{R}_+$ automatically translates into a partitioning of $\mathbb{R}$. For parts of the proofs, however, it is more convenient to work with the following alternative (albeit equivalent) definition.

Definition E.1. *Define $\mathcal{Q}_1 = -\mathcal{R}_K$, $\mathcal{Q}_2 = -\mathcal{R}_{K-1}, \ldots, \mathcal{Q}_K = -\mathcal{R}_1$, $\mathcal{Q}_{K+k} = \mathcal{R}_k$, $k \in [K]$, and $\widetilde{\mu} = (-\mu_K, \ldots, \mu_1, \mu_1, \ldots, \mu_K)^\top$. Then an equivalent definition of the quantization map is given by $z \mapsto Q(z) = \sum_{k=1}^{2K} \widetilde{\mu}_k I(z \in \mathcal{Q}_k)$. Likewise, we define $\widetilde{\mathbf{t}} = (-t_K, -t_{K-1}, \ldots, t_0, t_1, \ldots, t_{K-1}, t_K)^\top$.*

Expression for λ. Recall that $\lambda = \lambda_{b,\sigma} = \lambda_{b,\sigma}(\mathbf{t}, \boldsymbol{\mu})$ is defined by $\lambda = \mathbf{E}[g\,\theta(g)]$, $g \sim N(0,1)$, where the map θ is in turn defined by the relation $\mathbf{E}[y_1|a_1] = \theta(z_1)$ (here and below $z_i = \langle a_i, x_u^* \rangle$,

$i \in [m]$). We have

$$\begin{aligned}
\mathbf{E}[y_1|a_1] &= \sum_{k=1}^{2^b} \widetilde{\mu}_k \, \mathbf{P}(y_1 \in \mathcal{Q}_k) \\
&= \sum_{k=1}^{2^b} \widetilde{\mu}_k \, \mathbf{P}(z_1 + \sigma\varepsilon_1 \in \mathcal{Q}_k) \\
&= \sum_{k=1}^{2^b} \widetilde{\mu}_k \, \mathbf{P}(z_1 + \sigma\varepsilon_1 \in (\widetilde{t}_k, \widetilde{t}_{k+1})) \\
&= \sum_{k=1}^{2^b} \widetilde{\mu}_k \left\{ \Phi((\widetilde{t}_{k+1} - z_1)/\sigma) - \Phi((\widetilde{t}_k - z_1)/\sigma) \right\}.
\end{aligned}$$

We conclude that the map θ is defined by

$$\theta(z) = \sum_{k=1}^{2^b} \widetilde{\mu}_k \left\{ \Phi((\widetilde{t}_{k+1} - z)/\sigma) - \Phi((\widetilde{t}_k - z)/\sigma) \right\}.$$

Next we invoke Lemma E.1 which yields $\lambda = \mathbf{E}[z\theta(z)] = \mathbf{E}[\theta'(z)]$. We have

$$\theta'(z) = \sum_{k=1}^{2^b} \widetilde{\mu}_k \left\{ \frac{1}{\sigma}\phi((z - \widetilde{t}_k)/\sigma) - \frac{1}{\sigma}\phi((z - \widetilde{t}_{k+1})/\sigma) \right\}.$$

With the help of Lemma E.2, we compute

$$\begin{aligned}
\mathbf{E}[\theta'(z)] &= \sum_{k=1}^{2^b} \widetilde{\mu}_k \int_{\mathbb{R}} \left\{ \frac{1}{\sigma}\phi((z - \widetilde{t}_k)/\sigma) - \frac{1}{\sigma}\phi((z - \widetilde{t}_{k+1})/\sigma) \right\} \phi(z) \, dz. \\
&= \sum_{k=1}^{2^b} \widetilde{\mu}_k \frac{1}{\sqrt{1+\sigma^2}} \left\{ \phi\left(\frac{\widetilde{t}_k}{\sqrt{1+\sigma^2}} \right) - \phi\left(\frac{\widetilde{t}_{k+1}}{\sqrt{1+\sigma^2}} \right) \right\}.
\end{aligned}$$

Applying Lemma E.3, the last expression can be rewritten as follows:

$$\begin{aligned}
&\sum_{k=1}^{2^b} \widetilde{\mu}_k \frac{1}{\sqrt{1+\sigma^2}} \left\{ \phi\left(\frac{\widetilde{t}_k}{\sqrt{1+\sigma^2}} \right) - \phi\left(\frac{\widetilde{t}_{k+1}}{\sqrt{1+\sigma^2}} \right) \right\} \\
&= \sum_{k=1}^{2^b} \widetilde{\mu}_k \frac{\mathbf{E}[\widetilde{g}|\widetilde{g} \in (\widetilde{t}_k, \widetilde{t}_{k+1})]}{1+\sigma^2} \left\{ \Phi\left(\widetilde{t}_{k+1}/\sqrt{1+\sigma^2} \right) - \Phi\left(\widetilde{t}_k/\sqrt{1+\sigma^2} \right) \right\}, \\
&\qquad\qquad \text{where } \ \widetilde{g} \sim N(0, 1+\sigma^2) \\
&= \frac{1}{1+\sigma^2} \sum_{k=1}^{2^b} \widetilde{\mu}_k \, \mathbf{E}[\widetilde{g}|\widetilde{g} \in \mathcal{Q}_k] \, \mathbf{P}(\widetilde{g} \in \mathcal{Q}_k) \\
&= \frac{1}{1+\sigma^2} \sum_{k=1}^{K} \mu_k \, \mathbf{E}[\widetilde{g}|\widetilde{g} \in \mathcal{R}_k] \, \mathbf{P}(|\widetilde{g}| \in \mathcal{R}_k) \\
&= \frac{1}{1+\sigma^2} \left\langle \boldsymbol{\alpha}(\mathbf{t}), \boldsymbol{E}(\mathbf{t}) \odot \boldsymbol{\mu} \right\rangle,
\end{aligned}$$

where the penultimate line follows from the symmetry of the Gaussian distribution around zero; at this point, we convert the partitioning of $\mathbb{R}$ into $\{\mathcal{Q}_k\}_{k=1}^{2K}$ back to the partitioning of $\mathbb{R}_+$ into $\{\mathcal{R}_k\}_{k=1}^{K}$ (cf. the remark preceding Definition E.1). The last line follows by comparison with the definitions in Lemma 1.

Expression for Ψ. Recall that $\Psi = \Psi_{b,\sigma} = \Psi_{b,\sigma}(\mathbf{t}, \boldsymbol{\mu})$ is given by $\sqrt{\mathbf{E}[y_1^2]}$. Note that the random variable y_1^2 takes values in $\{\mu_k^2\}_{k=1}^K$ with $\mathbf{P}(y_1^2 = \mu_k^2) = \alpha_k(\mathbf{t})$, $k \in [K]$. Accordingly,

$$\sqrt{\mathbf{E}[y_1^2]} = \sqrt{\sum_{k=1}^{K} \alpha_k(\mathbf{t})\mu_k^2} = \sqrt{\langle \boldsymbol{\alpha}(\mathbf{t}), \boldsymbol{\mu} \odot \boldsymbol{\mu} \rangle}.$$

F Proof of Theorem 3

Consider the optimization problem

$$\min_{\mathbf{t},\boldsymbol{\mu}} \Omega_{b,\sigma}(\mathbf{t}, \boldsymbol{\mu}) = \min_{\mathbf{t},\boldsymbol{\mu}} \frac{\Psi_{b,\sigma}(\mathbf{t}, \boldsymbol{\mu})}{\lambda_{b,\sigma}(\mathbf{t}, \boldsymbol{\mu})}.$$

By Lemma 1, the above minimization problem is equivalent to

$$\min_{\mathbf{t},\boldsymbol{\mu}} R(\mathbf{t}, \boldsymbol{\mu}), \quad R(\mathbf{t}, \boldsymbol{\mu}) = \frac{\sqrt{\langle \boldsymbol{\alpha}(\mathbf{t}), \boldsymbol{\mu} \odot \boldsymbol{\mu} \rangle}}{\langle \boldsymbol{\alpha}(\mathbf{t}), \boldsymbol{E}(\mathbf{t}) \odot \boldsymbol{\mu} \rangle}, \tag{36}$$

where the term $\sigma^2 + 1$ in $\lambda_{b,\sigma}$ has been dropped as it does not depend on $\mathbf{t}$ or $\boldsymbol{\mu}$. We start by claiming that

$$R(\mathbf{t}, \boldsymbol{\mu}) \geq \frac{1}{\sqrt{\langle \boldsymbol{\alpha}(\mathbf{t}), \boldsymbol{E}(\mathbf{t}) \odot \boldsymbol{E}(\mathbf{t}) \rangle}} \tag{37}$$

for all $\boldsymbol{\mu}$ with distinct, non-zero entries. The above lower bound is attained by choosing $\boldsymbol{\mu}$ proportional to $\boldsymbol{E}(\mathbf{t})$ (note that the minimizing $\boldsymbol{\mu}$ is only defined up to a positive constant as $R(\mathbf{t}, c\boldsymbol{\mu}) = R(\mathbf{t}, \boldsymbol{\mu})$ for all $c > 0$). Inequality (37) follows from the Cauchy-Schwarz inequality. Denote by $\boldsymbol{A}(\mathbf{t})$ the diagonal matrix whose diagonal is given by the entries of $\boldsymbol{\alpha}(\mathbf{t})$. We then have

$$\begin{aligned}\langle \boldsymbol{\alpha}(\mathbf{t}), \boldsymbol{E}(\mathbf{t}) \odot \boldsymbol{\mu} \rangle &= \left\langle \boldsymbol{A}^{1/2}(\mathbf{t})\boldsymbol{E}(\mathbf{t}), \boldsymbol{A}^{1/2}(\mathbf{t})\boldsymbol{\mu} \right\rangle \\ &\leq \sqrt{\left\langle \boldsymbol{A}^{1/2}(\mathbf{t})\boldsymbol{E}(\mathbf{t}), \boldsymbol{A}^{1/2}(\mathbf{t})\boldsymbol{E}(\mathbf{t}) \right\rangle} \sqrt{\left\langle \boldsymbol{A}^{1/2}(\mathbf{t})\boldsymbol{\mu}, \boldsymbol{A}^{1/2}(\mathbf{t})\boldsymbol{\mu} \right\rangle}\end{aligned}$$

with equality holding if and only if

$$\boldsymbol{A}^{1/2}(\mathbf{t})\boldsymbol{E}(\mathbf{t}) = c\boldsymbol{A}^{1/2}(\mathbf{t})\boldsymbol{\mu} \;\Leftrightarrow\; \boldsymbol{E}(\mathbf{t}) = c\boldsymbol{\mu},$$

for some $c > 0$, where the above $\Leftrightarrow$ follows from the fact that the entries of $\mathbf{t}$ are required to be distinct so that the matrix $\boldsymbol{A}^{1/2}$ is regular. We conclude that

$$\min_{\mathbf{t},\boldsymbol{\mu}} R(\mathbf{t}, \boldsymbol{\mu}) = \min_{\mathbf{t}} R(\mathbf{t}, \boldsymbol{E}(\mathbf{t})) = \min_{\mathbf{t}} \frac{1}{\sqrt{\langle \boldsymbol{\alpha}(\mathbf{t}), \boldsymbol{E}(\mathbf{t}) \odot \boldsymbol{E}(\mathbf{t}) \rangle}}. \tag{38}$$

We will now show that the above minimization problem in $\mathbf{t}$ is equivalent to the b-bit Lloyd-Max quantization problem (16) of a random variable $h \sim N(0, 1 + \sigma^2)$, which we re-state here for convenience:

$$\min_{\mathbf{t},\boldsymbol{\mu}} \mathbf{E}[\{h - Q(h; \mathbf{t}, \boldsymbol{\mu})\}^2] = \min_{\mathbf{t},\boldsymbol{\mu}} \mathbf{E}[\{h - \operatorname{sign}(h) \textstyle\sum_{k=1}^{K} \mu_k I(|h| \in \mathcal{R}_k(\mathbf{t}))\}^2] \tag{39}$$

For the above problem, it is not hard to see that for any fixed choice of $\mathbf{t}$, the minimizing $\boldsymbol{\mu}^*(\mathbf{t})$ is given by $\mu_k^*(\mathbf{t}) = \mathbf{E}[h|h \in \mathcal{R}_k(\mathbf{t})] = E_k(\mathbf{t})$, $k \in [K]$, where we recall that $E_k(\mathbf{t})$ is the k-th component of $\boldsymbol{E}(\mathbf{t})$ as appearing above. To finish the proof of the first part of the Theorem 1, it thus remains to show that after substituting $\boldsymbol{\mu}^*(\mathbf{t})$ back into (39), the resulting minimization problem in $\mathbf{t}$ is equivalent to (38). We have

$$\begin{aligned}
&\min_{\mathbf{t}} \mathbf{E}\left[\left\{h - \operatorname{sign}(h)\sum_{k=1}^{K} I(|h| \in \mathcal{R}_k(\mathbf{t})\,)\,\mathbf{E}[h|h \in \mathcal{R}_k(\mathbf{t})]\right\}^2\right] \\
&= 2\min_{\mathbf{t}} \mathbf{E}\left[\sum_{k=1}^{K} I(h \in \mathcal{R}_k(\mathbf{t}))(h - \mathbf{E}[h|h \in \mathcal{R}_k(\mathbf{t})])^2\right] \\
&= 2\min_{\mathbf{t}} \mathbf{E}\left[\sum_{k=1}^{K} I(h \in \mathcal{R}_k(\mathbf{t}))\left\{h^2 - 2h\ \mathbf{E}[h|h \in \mathcal{R}_k(\mathbf{t})] + \mathbf{E}[h|h \in \mathcal{R}_k(\mathbf{t})]^2\right\}\right] \\
&= \mathbf{E}[h^2] + 2\min_{\mathbf{t}}\Bigg\{ -2\sum_{k=1}^{K} \mathbf{E}[h|h \in \mathcal{R}_k(\mathbf{t})]\,\mathbf{E}[I(h \in \mathcal{R}_k(\mathbf{t}))h] + \\
&\qquad\qquad + \sum_{k=1}^{K} \mathbf{P}(h \in \mathcal{R}_k(\mathbf{t}))\,\mathbf{E}[h|h \in \mathcal{R}_k(\mathbf{t})]^2\Bigg\} \\
&= 1 + \min_{\mathbf{t}} -\sum_{k=1}^{K} \mathbf{E}[h|h \in \mathcal{R}_k(\mathbf{t})]^2\,\mathbf{P}(|h| \in \mathcal{R}_k(\mathbf{t})) \\
&= 1 + \min_{\mathbf{t}} -\langle \boldsymbol{E}(\mathbf{t}) \odot \boldsymbol{E}(\mathbf{t}), \boldsymbol{\alpha}(\mathbf{t})\rangle\,,
\end{aligned}$$

which establishes the equivalence to (38) as claimed.

We now prove the second part of the theorem. Denote by $\mathbf{t}_0^*$ the Lloyd-Max optimal thresholds for $\sigma = 0$, i.e., for a $N(0,1)$ variable. Clearly, $\mathbf{t}^* = \mathbf{t}_\sigma^* = \sqrt{1+\sigma^2}\mathbf{t}_0^*$ for any $\sigma > 0$. Evaluating $\Omega_{b,\sigma}(\mathbf{t}^*, \boldsymbol{\mu}^*)$, we obtain in view of Lemma 1 and (36),(38)

$$\begin{aligned}
\Omega_{b,\sigma}(\mathbf{t}^*, \boldsymbol{\mu}^*) &= \frac{1+\sigma^2}{\sqrt{\langle \boldsymbol{\alpha}(\mathbf{t}^*), \boldsymbol{E}(\mathbf{t}^*) \odot \boldsymbol{E}(\mathbf{t}^*)\rangle}} \\
&= \frac{1+\sigma^2}{\sqrt{\left\langle \boldsymbol{\alpha}(\mathbf{t}_0^*\sqrt{1+\sigma^2}), \boldsymbol{E}(\mathbf{t}_0^*\sqrt{1+\sigma^2}) \odot \boldsymbol{E}(\mathbf{t}_0^*\sqrt{1+\sigma^2})\right\rangle}}
\end{aligned}$$

Evaluating the expression in the denominator, we obtain that

$$\alpha_k(\mathbf{t}_0^*\sqrt{1+\sigma^2}) = \mathbf{P}(|\widetilde{g}| \in \mathcal{R}_k(\mathbf{t}_0^*\sqrt{1+\sigma^2})) = \mathbf{P}(|g| \in \mathcal{R}_k(\mathbf{t}_0^*)),\ k \in [K],$$

where $\widetilde{g} \sim N(0, 1+\sigma^2)$, $g \sim N(0,1)$. Moreover, with the help of Lemma E.3

$$\begin{aligned}
\boldsymbol{E}(\mathbf{t}_0^*\sqrt{1+\sigma^2}) &= \Big(\mathbf{E}[\widetilde{g}|\widetilde{g} \in \mathcal{R}_k(\mathbf{t}_0^*\sqrt{1+\sigma^2})]\Big)_{k=1}^K \\
&= \sqrt{1+\sigma^2}\Big(\mathbf{E}[g|g \in \mathcal{R}_k(\mathbf{t}_0^*)]\Big)_{k=1}^K.
\end{aligned}$$

Putting together the pieces, we obtain that

$$\Omega_{b,\sigma}(\mathbf{t}^*, \boldsymbol{\mu}^*) = \frac{1+\sigma^2}{\sqrt{\langle \boldsymbol{\alpha}_0(\mathbf{t}_0^*), (1+\sigma^2)\boldsymbol{E}_0(\mathbf{t}_0^*) \odot \boldsymbol{E}_0(\mathbf{t}_0^*)\rangle}} = \frac{\sqrt{1+\sigma^2}}{\sqrt{\lambda_{b,0}(\mathbf{t}_0^*, \boldsymbol{\mu}_0^*)}}$$

where the $\boldsymbol{\alpha}_0(\mathbf{t})$, $\boldsymbol{E}_0(\mathbf{t})$ and $\lambda_{b,0}(\mathbf{t},\boldsymbol{\mu})$ refer to the definitions of $\boldsymbol{\alpha}(\mathbf{t}), \boldsymbol{E}(\boldsymbol{t}), \lambda_{b,\sigma}(\mathbf{t},\boldsymbol{\mu})$ for $\sigma = 0$. This completes the proof.

G Proof of Proposition 1

In the sequel, we derive tail bounds of the form

$$\begin{aligned}\mathbf{P}(\widehat{\psi} \geq (1+\varepsilon)\psi^*) &\leq \exp(-cm\varepsilon^2),\\ \mathbf{P}(\widehat{\psi} \leq (1-\varepsilon)\psi^*) &\leq \exp(-2cm\varepsilon^2).\end{aligned}$$

for $\varepsilon \in (0,1)$ and $c = 2\{\phi'(t/\psi^*)\}^2$. This implies that the probability of the event $\left\{\left|\frac{\widehat{\psi}}{\psi^*} - 1\right| > \varepsilon\right\}$ is upper bounded by $2\exp(-cm\varepsilon^2)$.

1) Upper tail

$$\begin{aligned}\mathbf{P}(\widehat{\psi} \geq (1+\varepsilon)\psi^*) &= \mathbf{P}\left(\frac{t_1}{\Phi^{-1}\left(\frac{1}{2}(1+\frac{m_1}{m})\right)} \geq (1+\varepsilon)\psi^*\right)\\ &= \mathbf{P}\left(\frac{m_1}{m} \leq 2\Phi\left(\frac{t_1}{(1+\varepsilon)\psi^*}\right) - 1\right)\\ &= \mathbf{P}\left(\frac{m_1}{m} - \mathbf{E}\left[\frac{m_1}{m}\right] \leq 2\left\{\Phi\left(\frac{t_1}{(1+\varepsilon)\psi^*}\right) - \Phi\left(\frac{t_1}{\psi^*}\right)\right\}\right)\end{aligned}$$

We have

$$\begin{aligned}\Phi\left(\frac{t_1}{(1+\varepsilon)\psi^*}\right) - \Phi\left(\frac{t_1}{\psi^*}\right) &= -\int_{t_1/(\psi^*(1+\varepsilon))}^{t_1/\psi^*} \phi(u)\, du\\ &\leq -\phi(t_1/\psi^*)\frac{t_1}{\psi^*}\frac{\varepsilon}{\varepsilon+1}\\ &\leq -\phi(t_1/\psi^*)\frac{t_1}{\psi^*}\frac{\varepsilon}{2} = \phi'(t_1/\psi^*)\frac{\varepsilon}{2}.\end{aligned}$$

for $\varepsilon \in (0,1)$. Thus

$$\mathbf{P}(\widehat{\psi} \geq (1+\varepsilon)\psi^*) \leq \mathbf{P}\left(\frac{m_1}{m} - \mathbf{E}\left[\frac{m_1}{m}\right] \leq \varepsilon\phi'(t_1/\psi^*)\right)$$

2) Lower tail

Similarly, we obtain that

$$\mathbf{P}(\widehat{\psi} \leq (1-\varepsilon)\psi^*) \leq \mathbf{P}\left(\frac{m_1}{m} - \mathbf{E}\left[\frac{m_1}{m}\right] \geq 2\left\{\Phi\left(\frac{t_1}{(1-\varepsilon)\psi^*}\right) - \Phi\left(\frac{t_1}{\psi^*}\right)\right\}\right)$$

We have

$$\Phi\left(\frac{t_1}{(1-\varepsilon)\psi^*}\right) - \Phi\left(\frac{t_1}{\psi^*}\right) = \int_{t_1/\psi^*}^{t_1/(\psi^*(1-\varepsilon))} \phi(u)\, du \qquad \geq \phi(t_1/\psi^*)\frac{t_1}{\psi^*}\frac{\varepsilon}{1-\varepsilon} \geq -\phi'(t_1/\psi^*)\varepsilon.$$

Thus,

$$\mathbf{P}(\widehat{\psi} \leq (1-\varepsilon)\psi^*) \leq \mathbf{P}\left(\frac{m_1}{m} - \mathbf{E}\left[\frac{m_1}{m}\right] \leq 2\varepsilon(-\phi'(t_1/\psi^*))\right).$$

Note that m_1 is a Binomial random variable. Applying Hoeffding's inequality to **1)** and **2)**, we obtain that

$$\mathbf{P}(\widehat{\psi} \geq (1+\varepsilon)\psi^*) \leq \exp\left(-2m\varepsilon^2\{\phi'(t_1/\psi^*)\}^2\right)$$
$$\mathbf{P}(\widehat{\psi} \leq (1-\varepsilon)\psi^*) \leq \exp\left(-4m\varepsilon^2\{\phi'(t_1/\psi^*)\}^2\right).$$

which proves the claim made above.

H Equivalence of $\widehat{x}$ and $\widehat{x}_\lambda$

We here derive that $\widehat{x}_\lambda/\|\widehat{x}_\lambda\|_2 = \widehat{x}$ if $\|\widehat{x}_\lambda\|_2 > 0$ and $\widehat{x}_\lambda = \widehat{x} = 0$ otherwise, with $\widehat{x}$ and $\widehat{x}_\lambda$ defined in (4) and (5), respectively, provided that $\mathcal{K}$ is a cone (i.e., $\alpha\mathcal{K} = \mathcal{K}$ for all $\alpha > 0$).

Let us consider the minimization problem that defines $\widehat{x}_\lambda$. Since $\mathcal{K}$ is a cone, we have for any $z \in \mathbb{R}^n$

$$\min_{x\in\lambda\mathcal{K}} \|x-z\|_2^2 = \min_{x\in\mathcal{K}} \|x-z\|_2^2 = \min_{\substack{u\in\mathcal{K}_1 \\ \alpha\geq 0}} \|\alpha u - z\|_2^2, \quad \text{where } \mathcal{K}_1 := \mathcal{K}\cap\mathbb{S}^{n-1}. \tag{40}$$

The latter optimization problem is equivalent to

$$\min_{\substack{u\in\mathcal{K}_1 \\ \alpha\geq 0}} \frac{1}{2}\alpha^2 - \alpha\langle u, z\rangle. \tag{41}$$

Note that for any fixed u, the minimizing α is given by $\alpha^*(u) = \max\{\langle u, z\rangle, 0\}$.

Case 1:
If $\max_{u\in\mathcal{K}_1} \alpha^*(u) = 0$, then 0 is the unique minimizer of (40). This immediately implies that 0 is also a minimizer of $\min_{x\in\mathcal{C}} -\langle x, z\rangle$, where we recall that $\mathcal{C} = \mathcal{K}\cap B_2^n$.

Case 2:
Otherwise, we then must have $\min_{z\in\mathcal{C}} -\langle x, z\rangle = \min_{x\in\mathcal{K}\cap\mathbb{S}^{n-1}} -\langle x, z\rangle = \min_{x\in\mathcal{K}_1} -\langle x, z\rangle$. On the other hand, substituting the expression for $\alpha^*(u)$ back into (41), we obtain the optimization problem $\min_{u\in\mathcal{K}_1} -\frac{1}{2}\langle u, z\rangle^2$, which is equivalent to the previous one. Hence $\widehat{x}_\lambda$ and $\widehat{x}$ only differ by a scalar multiple.

I Computation of $\widehat{x}$ for $\mathcal{K} = B_0(s;n)$

Lemma I.1. *Consider the optimization problem in* (4) *with* $\mathcal{K} = B_0(s;n)$, *i.e.*

$$\widehat{x} \in \operatorname*{argmin}_{x\in B_0(s;n)\cap B_2^n} -\langle \eta, x\rangle$$

Suppose that η *has distinct entries so that* $|\eta_{(1)}| > \ldots > |\eta_{(n)}|$, *where* $|\eta_{(k)}|$ *denotes the k-th largest value of* $\{|\eta_j|,\ j\in[n]\}$, $k\in[n]$. *It then holds that* $\widehat{x} = \widetilde{x}/\|\widetilde{x}\|_2$, *where* $\widetilde{x}_j = \eta_j I(|\eta_j| \geq |\eta_{(s)}|),\ j\in[n]$.

We note that if the $\{a_i\}_{i=1}^n$ are Gaussian, η has distinct entries with probability one.

Proof. Let $\emptyset \neq S \subseteq \{1,\ldots,n\}$. Then for any unit vector x supported on S, $\langle \eta, x\rangle \leq \|\eta_S\|_2$ which is attained by setting $x_S = \eta_S/\|\eta_S\|_2$. Consequently,

$$\min_{x\in B_0(s;n)\cap B_2^n} -\langle \eta, x\rangle = \min_{S:|S|\leq s} -\|\eta_S\|_2.$$

The optimization problem on the right hand side can be solved by finding the index set of the s largest component (in absolute magnitude) in η. This yields the claim. □

J Gaussian widths for the examples in §3.3

1) *Sparsity.*

To make this paper self-contained and since examples 2) and 3) depend on this result as well, we repeat the argument from the proof of Lemma 2.3 in [33] at this point, and provide explicit constants. For $1 \leq r \leq n$, note that $B_0(r;n) \cap \mathbb{S}^{n-1}$ is the union over the unit spheres of all $\binom{n}{r}$ subspaces of dimension r in $\mathbb{R}^n$. Accordingly, we obtain that

$$\mathbf{E}\left[\sup_{v \in B_0(r;n) \cap \mathbb{S}^{n-1}} |\langle g, v\rangle|\right] = \mathbf{E}\left[\sup_{\substack{S \subseteq [n], |S|=r \\ v: \|v\|_2 = 1, v_{S^c}=0}} |\langle g, v\rangle|\right] \leq \mathbf{E}\left[\max_{S \subseteq [n], |S|=r} \|g_S\|_2\right].$$

The latter expectation can be controlled by integrating and using a tail bound combined with the union bound:

$$\begin{aligned}\mathbf{E}\left[\sup_{S \subseteq [n], |S|=r} \|g_S\|_2\right] &= \int_0^\infty \mathbf{P}\left(\max_{S \subseteq [n], |S|=r} \|g_S\|_2 > u\right) du \\ &= \int_0^a \mathbf{P}\left(\max_{S \subseteq [n], |S|=r} \|g_S\|_2 > u\right) du + \int_a^\infty \mathbf{P}\left(\max_{S \subseteq [n], |S|=r} \|g_S\|_2 > u\right) du,\end{aligned}$$

for any $a > 0$. Choosing $a = \sqrt{r} + \sqrt{2\log\binom{n}{r}}$, we obtain that

$$\begin{aligned}\mathbf{E}\left[\sup_{S \subseteq [n], |S|=r} \|g_S\|_2\right] &\leq \sqrt{r} + \sqrt{2\log\binom{n}{r}} + \int_0^\infty \mathbf{P}\left(\max_{S \subseteq [n], |S|=r} \|g_S\|_2 > a + u\right) du \\ &\leq \sqrt{r} + \sqrt{2\log\binom{n}{r}} + \int_0^\infty \binom{n}{r} \mathbf{P}\left(\|(g_1, \ldots, g_r)^\top\|_2 > a + u\right) du\end{aligned}$$

by applying the union bound. From $\mathbf{E}[\|(g_1, \ldots, g_r)^\top\|_2] \leq \sqrt{r}$ and the well-known tail bound $\mathbf{P}(\|(g_1, \ldots, g_r)^\top\|_2 - \mathbf{E}[\|(g_1, \ldots, g_r)^\top\|_2] > t) \leq \exp(-t^2/2)$ for any $t \geq 0$ (cf. Proposition 5.34 in [42]), we obtain that

$$\int_0^\infty \binom{n}{r} \mathbf{P}\left(\|(g_1, \ldots, g_r)^\top\|_2 > a + u\right) du \leq \int_0^\infty \exp(-u^2/2)\, du = \sqrt{\pi/2}.$$

Using that $\binom{n}{r} \leq \exp(r\log(en/r))$, putting together the pieces and setting $r = 2s$, the derivation is complete.

2) *Fused Sparsity.*

Let us start by fixing some notation. For $k \in [n]$, denote by $\mathcal{P}(k;n)$ the set of all collection of subsets of $[n]$ defining a partitioning into k blocks, i.e.,

$$\begin{aligned}\mathcal{P}(k;n) = \Big\{ \{B_\ell\}_{\ell=1}^k : \exists \{a_\ell\}_{\ell=1}^k \subset \mathbb{N},\ 1 = a_1 \leq \ldots \leq a_k \leq n \\ \text{s.t. } B_\ell = \{j \in [n] : a_\ell \leq j \leq a_{\ell+1} - 1\},\ \ell = 1, \ldots, k-1, \quad B_k = \{a_k, \ldots, n\}.\Big\}\end{aligned}$$

Moreover, for $B \subseteq [n]$ let $\mathbf{1}_B \in \mathbb{R}^n$ denote the indicator vector of the index set B, i.e., $(\mathbf{1}_B)_j = 1$ if $j \in B$ and $(\mathbf{1}_B)_j = 0$ otherwise.

Note first that $\Delta(\mathcal{C}) \subset \mathrm{PC}(2s;n)$. We hence have

$$\overline{\Delta}(\mathcal{C}) \subset \mathrm{PC}(2s;n) \cap \mathbb{S}^{n-1} = \left\{ x \in \mathbb{R}^n : \sum_{\ell=1}^{2s} \mathbf{1}_{B_\ell} c_\ell,\ c_\ell \in \mathbb{R},\ \{B_\ell\}_{\ell=1}^{2s} \in \mathcal{P}(2s;n),\ \|x\|_2 = 1 \right\}$$
$$= \left\{ x \in \mathbb{R}^n : x = \sum_{\ell=1}^{2s} \frac{\mathbf{1}_{B_\ell}}{\sqrt{|B_\ell|}} \theta_\ell,\ \{B_\ell\}_{\ell=1}^{2s} \in \mathcal{P}(2s;n),\ \|\theta\|_2 = 1 \right\}.$$

Accordingly, we obtain that

$$w(\overline{\Delta}(\mathcal{C})) = \mathbf{E}\left[\sup_{z \in \overline{\Delta}(\mathcal{C})} |\langle g, z \rangle| \right] = \mathbf{E}\left[\sup_{\substack{\theta \in \mathbb{S}^{2s-1} \\ \{B_\ell\}_{\ell=1}^{2s} \in \mathcal{P}(2s;n)}} \left| \sum_{\ell=1}^{2s} \left\langle g, \mathbf{1}_{B_\ell} / \sqrt{|B_\ell|} \right\rangle \theta_\ell \right| \right]$$

For $\{B_\ell\}_{\ell=1}^{2s} \in \mathcal{P}(2s;n)$ arbitrary, consider the random vector $g_{\{B_\ell\}_{\ell=1}^{2s}} = \left(\left\langle g, \mathbf{1}_{B_\ell} / \sqrt{|B_\ell|} \right\rangle \right)_{\ell=1}^{2s}$. Observe that $g_{\{B_\ell\}_{\ell=1}^{2s}} \sim N(0, I_{2s})$. Therefore

$$\mathbf{E}\left[\sup_{\substack{\theta \in \mathbb{S}^{2s-1} \\ \{B_\ell\}_{\ell=1}^{2s} \in \mathcal{P}(2s;n)}} \left| \sum_{\ell=1}^{2s} \left\langle g, \mathbf{1}_{B_\ell} / \sqrt{|B_\ell|} \right\rangle \theta_\ell \right| \right] = \mathbf{E}\left[\sup_{\substack{\theta \in \mathbb{S}^{2s-1} \\ \{B_\ell\}_{\ell=1}^{2s} \in \mathcal{P}(2s;n)}} \left| \left\langle g_{\{B_\ell\}_{\ell=1}^{2s}}, \theta \right\rangle \right| \right]$$
$$= \mathbf{E}\left[\sup_{\{B_\ell\}_{\ell=1}^{2s} \in \mathcal{P}(2s;n)} \left\| g_{\{B_\ell\}_{\ell=1}^{2s}} \right\|_2 \right]$$

Noting that $|\mathcal{P}(2s;n)| \leq \binom{n-1}{2s-1}$, the last expectation can be bounded as in the derivation of 1) above.

3) *Group sparsity.*

We have $\Delta(\mathcal{C}) \subset B_{0,\mathcal{G}}(2s)$ and hence $w(\overline{\Delta}(\mathcal{C})) \leq w(B_{0,\mathcal{G}}(2s) \cap \mathbb{S}^{n-1})$. Analogously to 1), we have

$$w(B_{0,\mathcal{G}}(2s) \cap \mathbb{S}^{n-1}) = \mathbf{E}\left[\sup_{v \in B_{0,\mathcal{G}}(2s) \cap \mathbb{S}^{n-1}} |\langle g, v \rangle| \right] \leq \mathbf{E}\left[\sup_{S \in [L], |S| \leq 2s} \|g_{I(S)}\|_2 \right],\ I(S) := \cup_{\ell \in S} G_\ell.$$

Following along the lines of i), we obtain

$$w(B_{0,\mathcal{G}}(2s)) \leq \sqrt{2s \max_{\ell \in [L]} |G_l|} + \sqrt{4s \log(eL/2s)} + \sqrt{\pi/2},$$

which concludes the derivation.

4) *Low-rank matrices.*

We have

$$\overline{\Delta}(\mathcal{C}) \subset \{X \in \mathbb{R}^{n_1 \times n_2} : \mathrm{rank}(X) \leq 2r,\ \|X\|_F \leq 1\} \subset \{X \in \mathbb{R}^{n_1 \times n_2} : \|X\|_* \leq \sqrt{2r},\ \|X\|_F \leq 1\},$$

where $\|\cdot\|_*$ denotes the Schatten-one norm (the sum of the singular values). Let G be an n_1-by-n_2 random matrix whose entries are i.i.d. $N(0,1)$. From the containment above, we obtain that

$$w(\overline{\Delta}(\mathcal{C})) \leq \mathbf{E}\left[\sup_{X \in \mathbb{R}^{n_1 \times n_2} : \|X\|_F \leq 1, \|X\|_* \leq \sqrt{2r}} |\langle G, X \rangle| \right] \leq \sqrt{2r}\, \mathbf{E}[\|G\|],$$

by the duality of the Schatten-one norm and the spectral norm. Gordon's Theorem (Theorem 5.3.2 in [42]) states that $\mathbf{E}[\|G\|] \le \sqrt{n_1} + \sqrt{n_2}$, which completes the derivation.

5) ℓ_1-*ball constraint.*

We have

$$\overline{\Delta}(\mathcal{C}) = \left\{ \frac{x - x_u^*}{\|x - x_u^*\|_2} \ : \ x \in B_2^n \setminus \{x_u^*\}, \ \|x\|_1 \le \|x_u^*\|_1 \right\}.$$

Denote by S, $|S| = s$, the support of x_u^*. For any x with $\|x\|_1 \le \|x_u^*\|_1$, we have $\|x_{S^c}\|_1 \le \|x_S - (x_u^*)_S\|_1$ and thus

$$\|x - x_u^*\|_1 \le \|x_{S^c}\|_1 + \|x_S - (x_u^*)_S\|_1 \le 2\|x_S - (x_u^*)_S\|_1 \le 2\sqrt{s}\|x - x_u^*\|_2.$$

It follows that

$$\overline{\Delta}(\mathcal{C}) \subset \{v : \|v\|_2 = 1, \ \|v\|_1 \le 2\sqrt{s} \ \} \subset \sqrt{2}\{v : \|v\|_2 \le 1, \ \|v\|_1 \le \sqrt{2s}\} \subset 2\sqrt{2}\,\mathrm{conv}\{B_0(2s;n) \cap \mathbb{S}^{n-1}\},$$

where the last containment is proved in [32], Lemma 3.1 therein. Since $w(\mathrm{conv}\{B_0(2s;n) \cap \mathbb{S}^{n-1}\}) = w(B_0(2s;n) \cap \mathbb{S}^{n-1})$, the derivation is complete.

K Derivations for the paragraph "Beyond additive noise" in §3.5

We fix $\sigma = 0$ and the corresponding Lloyd-Max optimal choices $\mathbf{t} = \mathbf{t}_0^*$, $\boldsymbol{\mu} = \boldsymbol{\mu}_0^*$ so that $\mu_k = \mathbf{E}[g | g \in \mathcal{R}_k]$, $g \sim N(0,1)$ with $\mathcal{R}_k = \mathcal{R}_k(\mathbf{t}_0^*)$, $k \in [K]$.

Mechanism (I). In order to evaluate $\lambda = \lambda_{b,p}$, we first need to derive an expression for the corresponding map θ. Recalling Definition E.1, we have

$$\mathbf{E}[y_1 | a_1] = (1-p) \sum_{k=1}^{2^b} \widetilde{\mu}_k I(\langle a_1, x^* \rangle \in \mathcal{Q}_k) + p \frac{1}{2^b - 1} \sum_{k=1}^{2^b} \widetilde{\mu}_k I(\langle a_1, x^* \rangle \notin \mathcal{Q}_k)$$

and thus

$$\theta(z) = (1-p) \sum_{k=1}^{2^b} \widetilde{\mu}_k I(z \in \mathcal{Q}_k) + p \frac{1}{2^b - 1} \sum_{k=1}^{2^b} \widetilde{\mu}_k I(z \notin \mathcal{Q}_k).$$

This is immediate from the definition of Mechanism (I) which keeps the original bin with probability $1 - p$, and selects one of the remaining $2^b - 1$ bins uniformly at random with probability p.

It follows that for $g \sim N(0,1)$

$$\begin{aligned}
\lambda_{b,p} = \mathbf{E}[g\,\theta(g)] &= \sum_{k=1}^{2^b} \widetilde{\mu}_k \left\{ (1-p)\,\mathbf{E}\left[gI(g \in \mathcal{Q}_k)\right] + p\frac{1}{2^b-1}\,\mathbf{E}[gI(g \notin \mathcal{Q}_k)] \right\} \\
&= \sum_{k=1}^{2^b} \widetilde{\mu}_k \left\{ (1-p)\,\mathbf{E}\left[gI(g \in \mathcal{Q}_k)\right] + p\frac{1}{2^b-1}\,\mathbf{E}[g(1-I(g \in \mathcal{Q}_k))] \right\} \\
&= \sum_{k=1}^{2^b} \widetilde{\mu}_k \left\{ (1-p) - \frac{p}{2^b-1} \right\} \mathbf{E}\left[gI(g \in \mathcal{Q}_k)\right] \\
&= \sum_{k=1}^{K} \mathbf{P}(|g| \in \mathcal{R}_k)\,\mathbf{E}[g|g \in \mathcal{R}_k]^2 \left\{ (1-p) - \frac{p}{2^b-1} \right\} \\
&= \langle \boldsymbol{\alpha}_0(\mathbf{t}_0^*), \boldsymbol{E}_0(\mathbf{t}_0^*) \odot \boldsymbol{E}_0(\mathbf{t}_0^*) \rangle \left\{ (1-p) - \frac{p}{2^b-1} \right\} \\
&= \lambda_{b,0} \left\{ (1-p) - \frac{p}{2^b-1} \right\},
\end{aligned}$$

where $\boldsymbol{\alpha}_0(\mathbf{t}_0^*)$ and $\boldsymbol{E}_0(\mathbf{t}_0^*)$ are defined at the end of Appendix F. From the last expression we deduce the breakdown point $\bar{p}_b = 1 - 1/2^b$.

$\Psi_{b,p}$ can be evaluated by using the expression in Lemma 1. The only thing that changes under Mechanism (I) are the probabilities $\boldsymbol{\alpha}(\mathbf{t})$ which become

$$\alpha_k(\mathbf{t}) = \mathbf{P}(|g| \in \mathcal{R}_k(\mathbf{t}))(1-p) + \frac{p}{2^{b-1}} \sum_{l \neq k} \mathbf{P}(|g| \in \mathcal{R}_l(\mathbf{t})), \quad k \in [K].$$

Mechanism (II). Following the same route as for Mechanism (I), one derives

$$\theta(z) = (1-p) \sum_{k=1}^{2^b} \widetilde{\mu}_k I(z \in \mathcal{Q}_k) + p\left\{ -\mu_K I(z \geq 0) + \mu_K I(z < 0) \right\}$$

and accordingly for $g \sim N(0,1)$

$$\begin{aligned}
\lambda_{b,p} = \mathbf{E}[g\,\theta(g)] &= (1-p) \sum_{k=1}^{K} \mathbf{P}(|g| \in \mathcal{R}_k)\,\mathbf{E}[g|g \in \mathcal{R}_k]^2 - p\mu_K\,\mathbf{E}[g|g > 0] \\
&= (1-p)\lambda_{b,0} - p\mu_K\sqrt{2/\pi}
\end{aligned}$$

so that the breakdown points results as $\bar{p}_b = \lambda_{b,0}/(\lambda_{b,0} + \mu_K\sqrt{2/\pi})$. As for Mechanism (I), $\Psi_{b,p}$ is obtained by evaluating the changes in $\boldsymbol{\alpha}(\mathbf{t})$. We have

$$\begin{aligned}
\alpha_k(\mathbf{t}) &= (1-p)\,\mathbf{P}(|g| \in \mathcal{R}_k(\mathbf{t})), \quad k \in [K-1], \\
\alpha_K(\mathbf{t}) &= p \sum_{k=1}^{K-1} \mathbf{P}(|g| \in \mathcal{R}_k(\mathbf{t})) + \mathbf{P}(|g| \in \mathcal{R}_K(\mathbf{t})).
\end{aligned}$$

L Auxiliary results from the theory of empirical processes

The first Lemma follows from a result known as symmetrization in the theory of empirical processes, cf. Lemma 2.3.1 in [41].

Lemma L.1. *Let Γ be a set of measurable functions and $\{Z_i\}_{i=1}^m$ random variables and $\{\xi_i\}_{i=1}^m$ be i.i.d. Rademacher random variables, i.e., $\mathbf{P}(\xi_1 = 1) = \mathbf{P}(\xi_1 = -1) = 1/2$. Then:*

$$\mathbf{E}\left[\sup_{\gamma\in\Gamma}\left|\frac{1}{m}\sum_{i=1}^m(\gamma(Z_i) - \mathbf{E}[\gamma(Z_i)])\right|\right] \le 2\,\mathbf{E}\left[\sup_{\gamma\in\Gamma}\left|\frac{1}{m}\sum_{i=1}^m \xi_i\gamma(Z_i)\right|\right]$$

The next lemma is essentially derived in [26], §6.1. We here also provide a detailed proof for the sake of completeness.

Lemma L.2. *In the setting of the previous Lemma, we have*

$$\begin{aligned}&\mathbf{P}\left(\sup_{\gamma\in\Gamma}\left|\frac{1}{m}\sum_{i=1}^m(\gamma(Z_i) - \mathbf{E}[\gamma(Z_i)])\right| > 2\,\mathbf{E}\left[\sup_{\gamma\in\Gamma}\left|\frac{1}{m}\sum_{i=1}^m(\gamma(Z_i) - \mathbf{E}[\gamma(Z_i)])\right|\right] + u\right) \qquad (42)\\ &\le 4\,\mathbf{P}\left(\sup_{\gamma\in\Gamma}\left|\frac{1}{m}\sum_{i=1}^m \xi_i\gamma(Z_i)\right| > u/2\right).\end{aligned}$$

Proof. Let Z_i' be an independent copy of Z_i, $i \in [m]$. Observe that for any $a, t > 0$, we have that

$$\begin{aligned}&\mathbf{P}\left(\sup_{\gamma\in\Gamma}\left|\frac{1}{m}\sum_{i=1}^m(\gamma(Z_i) - \mathbf{E}[\gamma(Z_i)])\right| \le a\right)\mathbf{P}\left(\sup_{\gamma\in\Gamma}\left|\frac{1}{m}\sum_{i=1}^m(\gamma(Z_i') - \mathbf{E}[\gamma(Z_i')])\right| > a + t\right)\\ \le&\,\mathbf{P}\left(\sup_{\gamma\in\Gamma}\left|\frac{1}{m}\sum_{i=1}^m(\gamma(Z_i) - \mathbf{E}[\gamma(Z_i)]) - \frac{1}{m}\sum_{i=1}^m(\gamma(Z_i') - \mathbf{E}[\gamma(Z_i')])\right| > t\right)\end{aligned}$$

Choose $a = 2\,\mathbf{E}\left[\sup_{\gamma\in\Gamma}\left|\frac{1}{m}\sum_{i=1}^m(\gamma(Z_i) - \mathbf{E}[\gamma(Z_i)])\right|\right]$. Then by Markov's inequality, the first term in the left hand side of the above inequality is lower bounded by $1/2$ and thus

$$\begin{aligned}&\mathbf{P}\left(\sup_{\gamma\in\Gamma}\left|\frac{1}{m}\sum_{i=1}^m(\gamma(Z_i') - \mathbf{E}[\gamma(Z_i')])\right| > 2\,\mathbf{E}\left[\sup_{\gamma\in\Gamma}\left|\frac{1}{m}\sum_{i=1}^m(\gamma(Z_i) - \mathbf{E}[\gamma(Z_i)])\right|\right] + t\right)\\ &\le 2\,\mathbf{P}\left(\sup_{\gamma\in\Gamma}\left|\frac{1}{m}\sum_{i=1}^m(\gamma(Z_i) - \mathbf{E}[\gamma(Z_i)]) - \frac{1}{m}\sum_{i=1}^m(\gamma(Z_i') - \mathbf{E}[\gamma(Z_i')])\right| > t\right)\\ &\le 2\,\mathbf{P}\left(\sup_{\gamma\in\Gamma}\left|\frac{1}{m}\sum_{i=1}^m \xi_i(\gamma(Z_i) - \gamma(Z_i'))\right| > t\right)\\ &\le 2\,\mathbf{P}\left(\sup_{\gamma\in\Gamma}\left|\frac{1}{m}\sum_{i=1}^m \xi_i\gamma(Z_i)\right| + \sup_{\gamma\in\Gamma}\left|\frac{1}{m}\sum_{i=1}^m \xi_i\gamma(Z_i')\right| > t\right)\\ &\le 4\,\mathbf{P}\left(\sup_{\gamma\in\Gamma}\left|\frac{1}{m}\sum_{i=1}^m \xi_i\gamma(Z_i)\right| > t/2\right)\end{aligned}$$

□

The next Lemma is taken from [6], Theorem 5.8.

Lemma L.3. *Let $\mathcal{T}$ be a totally bounded subset of some metric space and let $(Z_t)_{t\in\mathcal{T}}$ be an almost surely continuous Gaussian process indexed by $\mathcal{T}$ such that $\sigma_Z^2 := \sup_t \mathbf{E}[Z_t^2] < \infty$. Then $\forall r > 0$*

$$\mathbf{P}\left(\sup_t Z_t - \mathbf{E}\left[\sup_t Z_t\right] \ge r\right) \le \exp(-r^2/(2\sigma_Z^2)).$$

The following result can be found in [36].

Lemma L.4. *Let* $\mathbf{Z} = (Z_i)_{i=1}^m$ *be a sub-Gaussian random vector whose components* $(Z_i)_{i=1}^m$ *are i.i.d. zero-mean and variance* σ_Z^2*. Then:*

$$\mathbf{P}(\|\mathbf{Z}\| > \mathbf{E}[\|\mathbf{Z}\|] + t) \leq \exp(-ct^2/\sigma_Z^2),$$

where $c > 0$ *is a universal constant depending only on the sub-Gaussian norm of* Z_1/σ_Z*.*

References

[1] A. Moshtaghpour and L. Jacques and V. Cambareri and K. Degraux and C. De Vleeschouwer. Consistent Basis Pursuit for Signal and Matrix Estimates in Quantized Compressed Sensing. *IEEE Signal Processing Letters*, 23:25–29, 2016.

[2] R. Adler and J. Taylor. *Random Fields and Geometry*. Springer, 2007.

[3] R. Bellman. On the approximation of curves by line segments using dynamic programming. *Communications of the ACM*, 4:284, 1961.

[4] P. Billingsley. *Probability and Measure*. Wiley, 1995.

[5] T. Blumensath and M. Davies. Iterative hard thresholding for compressed sensing. *Applied and Computational Harmonic Analysis*, 27:265–274, 2009.

[6] S. Boucheron, G. Lugosi, and P. Massart. *Concentration Inequalities: A Nonasymptotic Theory of Independence*. Oxford University Press, 2013.

[7] P. Boufounos and R. Baraniuk. 1-Bit Compressive Sensing. In *Information Science and Systems*, 2008.

[8] P. Boufounos, L. Jacques, F. Krahmer, and R. Saab. Quantization and Compressed Sensing. In *Compressed Sensing and its Applications*, pages 193–237. Springer.

[9] D. Brillinger. A generalized linear model with "Gaussian" regressor variables. *A Festschrift For Erich L. Lehmann*, pages 97–114, 1982.

[10] T. Cai, T. Liang, and A. Rakhlin. Geometric inference for general high-dimensional linear inverse problems. *The Annals of Statistics, to appear*, 2016.

[11] E. Candes and M. Davenport. How well can we estimate a sparse vector ? *Applied and Computational Harmonic Analysis*, 34:317–323, 2013.

[12] E. Candes and Y. Plan. Tight oracle bounds for low-rank matrix recovery from a minimal number of noisy measurements. *IEEE Transactions on Information Theory*, 57:2342–2359, 2011.

[13] V. Chandrasekaran, B. Recht, P. Parrilo, and A. Willsky. The convex geometry of linear inverse problems. *Foundations of Computational Mathematics*, 12:805–849, 2012.

[14] Compressive Sensing Resources. http://dsp.rice.edu/cs.

[15] R.-E. Fan, K.-W. Chang, C.-J. Hsieh, X.-R. Wang, and C.-J. Lin. LIBLINEAR: A library for large linear classification. *Journal of Machine Learning Research*, 9:1871–1874, 2008.

[16] R. Gallager. Course materials for 'Principles of Digital Communications I'. MIT OpenCourseWare `http://ocw.mit.edu/`, 2006.

[17] S. Gopi, P. Netrapalli, P. Jain, and A. Nori. One-bit Compressed Sensing: Provable Support and Vector Recovery. In *International Conference on Machine Learning (ICML)*, 2013.

[18] A. Gupta, B. Recht, and R. Nowak. Sample complexity for 1-bit compressed sensing and sparse classification. In *International Symposium on Information Theory (ISIT)*, pages 1553–1557, 2010.

[19] L. Jacques. Personal communication.

[20] L. Jacques. Error Decay of (almost) Consistent Signal Estimations from Quantized Gaussian Random Projections. *arxiv:1406.0022; to appear in IEEE Transactions on Information Theory*, 2016.

[21] L. Jacques, K. Degraux, and C. De Vleeschouwer. Quantized iterative hard thresholding: Bridging 1-bit and high-resolution quantized compressed sensing. arXiv:1305.1786, 2013.

[22] L. Jacques, J. Laska, P. Boufounos, and R. Baraniuk. Robust 1-bit Compressive Sensing via Binary Stable Embeddings of Sparse Vectors. *IEEE Transactions on Information Theory*, 59:2082–2102, 2013.

[23] J. Kieffer. Uniqueness of locally optimal quantizer for log-concave density and convex error weighting function. *IEEE Transactions on Information Theory*, 29:42–47, 1983.

[24] S. Kim, K. Koh, S. Boyd, and D. Gorinevsky. l_1 trend filtering. *SIAM Review*, 51:339–360, 2009.

[25] J. Laska and R. Baraniuk. Regime change: Bit-depth versus measurement-rate in compressive sensing. *IEEE Transactions on Signal Processing*, 60:3496–3505, 2012.

[26] M. Ledoux and M. Talagrand. *Probability in Banach Spaces: Isoperimetry and Processes*. Springer, 1991.

[27] P. Li. Binary and Multi-Bit Coding for Stable Random Projections. arXiv:1503.06876, 2015.

[28] P. Li. One scan 1-bit compressed sensing. Technical report, arXiv:1503.02346, 2015.

[29] S. Lloyd. Least Squares Quantization in PCM. *IEEE Transactions on Information Theory*, 28:129–137, 1982.

[30] K. Lounici, M. Pontil, A. Tsybakov, and S. van de Geer. Oracle inequalities and Optimal Inference under Group Sparsity. *The Annals of Statistics*, 39:2164–2204.

[31] J. Max. Quantizing for Minimum Distortion. *IRE Transactions on Information Theory*, 6:7–12, 1960.

[32] Y. Plan and R. Vershynin. One-bit compressed sensing by linear programming. *Communications on Pure and Applied Mathematics*, 66:1275–1297, 2013.

[33] Y. Plan and R. Vershynin. Robust 1-bit compressed sensing and sparse logistic regression: a convex programming approach. *IEEE Transactions on Information Theory*, 59:482–494, 2013.

[34] Y. Plan and R. Vershynin. The generalized Lasso with non-linear observations. arXiv:1502.04071; to appear in IEEE Transactions on Information Theory, 2015.

[35] Y. Plan, R. Vershynin, and E. Yudovina. High-dimensionl estimation with geometric constraints. arXiv:1404.7349, 2014.

[36] M. Rudelson and R. Vershynin. Hanson-Wright inequality and sub-gaussian concentration. *Electronic Communications in Probability*, 18:1–9, 2013.

[37] M. Slawski and P. Li. b-bit marginal regression. In *Advances in Neural Information Processing Systems (NIPS)*, pages 2053–2061, 2015.

[38] C. Thrampoulidis, E. Abbasi, and B. Hassabi. The LASSO with Non-linear Measurements is Equivalent to One with Linear Measurements. arXiv:1506.02181v1, 2015.

[39] R. Tibshirani. Regression shrinkage and variable selection via the lasso. *Journal of the Royal Statistical Society Series B*, 58:671–686, 1996.

[40] A. van der Vaart. *Asymptotic Statistics*. Cambridge University Press, 1998.

[41] A. van der Vaart and J. Wellner. *Weak Convergence and Empirical Processes*. Springer, 1996.

[42] R. Vershynin. *In: Compressed Sensing: Theory and Applications*, chapter 'Introduction to the non-asymptotic analysis of random matrices'. Cambridge University Press, 2012.

[43] M. Yuan and Y. Lin. Model selection and estimation in regression with grouped variables. *Journal of the Royal Statistical Society Series B*, 68:49–67, 2006.

[44] L. Zhang, J. Yi, and R. Jin. Efficient algorithms for robust one-bit compressive sensing. In *International Conference on Machine Learning (ICML)*, 2014.

[45] A. Zymnis, S. Boyd, and E. Candes. Compressed sensing with quantized measurements. *Signal Processing Letters*, 17:149–152, 2009.